Cosmological Reionization by Stellar Sources

Nickolay Y. Gnedin
CASA, University of Colorado, Boulder, CO 80309; e-mail: *gnedin@casa.colorado.edu*

ABSTRACT

I use cosmological simulations that incorporate a physically motivated approximation to three-dimensional radiative transfer that recovers correct asymptotic ionization front propagation speeds for some cosmologically relevant density distributions to investigate the process of the reionization of the universe by ionizing radiation from proto-galaxies. Reionization proceeds in three stages and occupies a large redshift range from $z \sim 15$ until $z \sim 5$. During the first, "pre-overlap" stage, H II regions gradually expand into the low density IGM, leaving behind neutral high density protrusions. During the second, "overlap" stage, that occurs in about 10% of the Hubble time, H II regions merge and the ionizing background rises by a large factor. During the third, "post-overlap" stage, remaining high density regions are being gradually ionized as the required ionizing photons are being produced.

Residual fluctuations in the ionizing background reach significant (more than 10%) levels for the Lyman-alpha forest absorption systems with column densities above $10^{14} - 10^{15}\,\mathrm{cm}^{-2}$ at $z = 3$ to 4.

Subject headings: cosmology: theory - cosmology: large-scale structure of universe - galaxies: formation - galaxies: intergalactic medium

1. Introduction

Existing ground based observations of the CMB on sub-degree angular scales suggest that the gas content of the universe was mostly neutral since recombination at $z \sim 1000$ until about $z \sim 100$ (Bond & Jaffe 1998; Griffiths, Barbosa, & Liddle 1998), because earlier reionization would have brought the last scattering surface to lower redshift, smoothing the intrinsic CMB anisotropy. At the same time we know that the universe is highly ionized since $z \approx 5$, from observations of the spectra of quasars with the highest redshifts (Giallongo et al. 1994; Williger et al. 1994; Songaila et al. 1999). This change of the ionization state of the universe from neutral to highly ionized is called *reionization*.

Recent years witnessed a surge in research on reionization along two separate directions. Semi-analytical methods attempt to describe the general features of the evolution of the intergalactic medium (IGM) based on simple, ad hoc assumptions about the star formation history

and the density distribution in the IGM. Most of the previous works utilized a simple clumping factor to account for the inhomogeneity of the IGM (Giroux & Shapiro 1996; Tegmark et al. 1997; Madau, Meiksin, & Rees 1997; Ciardi & Ferrara 1997; Haiman & Loeb 1997, 1998; Madau, Haardt, & Rees 1999; Valageas & Silk 1999; Chiu & Ostriker 1999). A recent work by Miralda-Escudé, Haehnelt, & Rees (1999) attempted to advance further by introducing the distribution of the ionized fraction as a function of gas density which, besides everything else, provides a specific model for various clumping factors. Despite using a very simple (perhaps even oversimplified) ansatz for the ionized fraction - density distribution, their main conclusions agree remarkably well with the results of the simulations presented in this paper, as I will show below.

The main advantage of the semi-analytical approaches is their simplicity and the ability to emphasize the key physical processes. Their main problem is a significant oversimplification and inability to follow the complex dynamical interactions between the dark matter, gas, stars, and the spatially inhomogeneous and time-variable radiation field. Semi-analytical models also miss the effects of correlation between the sources of ionizing radiation.

The latter deficiency was partly overcome recently by Ciardi et al. (1999) by combining a semi-analytical model with the N-body simulation, but the lack of dynamics will elude semi-analytical models forever.

Cosmological numerical simulations offer a totally different approach to modeling reionization (Ostriker & Gnedin 1996; Gnedin & Ostriker 1997). The main advantage of simulations over the semi-analytical models is that numerical simulations fully account for the dynamical evolution of the matter contents of the universe, thus avoiding the main over-simplification of the semi-analytical models. Their inherited limitation is a limited dynamic range, which for existing numerical codes implies an impossibility to achieve numerical convergence and therefore quantitatively accurate results. However, semi-analytical models are also imprecise since they adopt ad hoc assumptions and ansatzes, and thus we must accept the fact that we are not able to model reionization quantitatively accurate at this moment, and need to limit ourselves with understanding the key qualitative features of the physical processes that take place in the universe during that epochs.

Until recently, numerical simulations failed to include the three-dimensional radiative transfer in all its complexity, and therefore were not sufficient to properly model the dynamics of the ionizing radiation in the universe.

The understanding of the crucial role that the full 3D radiative transfer plays in the physics of reionization prompted the first attempts to develop numerical techniques to incorporate the radiative transfer into numerical simulations (Abel, Norman, & Madau 1999). While Abel et al. (1999) were able to develop an algorithm for an exact implementation of the full 3D radiative transfer into cosmological simulations, their algorithm is appropriate for following the radiation field in the case when there are only a few sources within the volume with the size of the mean free path of an ionizing photon. When the number of sources within this volume increases, as is

the case after all H II regions overlap, their method becomes prohibitively expensive.

In this paper I adopt a somewhat different approach. Rather than trying to model the radiative transfer exactly, I develop an *approximation*, which is sufficiently fast computationally. I validate the adopted approximation on a series of spherically symmetric test cases, for which the full radiative transfer can be solved exactly at a modest computational expense.

The physical basis for the adopted approximation, which I call a "Local Optical Depth" (LOD) approximation, is simple. In the optically thin regime it is straightforward to compute the spatially variable radiation field, as the problem reduces to simply collecting r^{-2} potential laws from all the sources in the simulation volume, which can be done fast with the P^{3}M algorithm. Thus, the key difficulty in following the full radiative transfer is to account for the optical depth, which can always be presented as a product of the cross-section, the gas density, and a characteristic scale. The LOD approximation adopts an ansatz for this characteristic scale as the scale over which the gas density changes significantly.

Having developed the means to follow approximately the 3D spatially inhomogeneous and time-dependent radiative transfer, I apply them to modeling the process of reionization of the universe in a specific scenario when sources of ionizing radiation are stars (grouped into proto-galaxies). There are two reasons why numerical simulations need to be restricted to this specific case at the moment.

The first reason is "scientific", as the mounting evidence suggests that quasars are unable to reionize the universe at $z > 5$ (Madau et al. 1999), and thus the universe is reionized by stars. The second reason is "pragmatic". A quasar radiation field may affect the properties of the IGM over a region several tens of megaparsecs in size. Thus, a numerical simulation capable of modeling the reionization by quasars would have to have a computational region of at least 100 Mpc in size. At the same time, in order to resolve the formation of the first structures and to follow the gas density with sufficient precision, it will need spatial resolution below 1 kpc and mass resolution better than about 10^6 solar mass in baryons. Such a dynamic range (10^5) is beyond the capabilities of the existing cosmological hydrodynamic codes.[1]

Thus, only a numerical simulation with the box size of several Mpc is practical at the moment. Fortunately, this box size is sufficient to model the reionization by stellar sources (on a semi-qualitative level), as will be shown below.

This paper is organized in a conventional way. In §2 I describe the simulations, §3 is devoted to the results, and §4 concludes with the discussion. Appendix describes the Local Optical Depth approximation and the tests of the method.

[1]With the exception of the Adaptive Mesh Refinement technique.

2. Simulations

Simulations reported in this paper were performed with the "Softened Lagrangian Hydrodynamics" (SLH-P^{3}M) code (Gnedin 1995, 1996; Gnedin & Bertschinger 1996). The following physical ingredients are currently included in the code:

Dark matter is followed using the adaptive P^{3}M algorithm.

Gas dynamics is followed on a quasi-Lagrangian deformable mesh using the SLH algorithm.

Star formation is included using the Schmidt law in resolution elements that sink below the numerical resolution of the code.

Atomic physics of hydrogen and helium plasma is followed exactly using a two-level implicit scheme.

Molecular hydrogen formation and destruction is followed exactly (including the radiative transfer effects) in the limit when the fraction of hydrogen in the molecular form is small and the self-shielding of H_2 is unimportant (in the latter case the approximate method of following the radiative transfer becomes exact). This is always the case in the simulation presented in this paper because the numerical resolution is not sufficient to resolve the formation of molecular clouds.

Radiative transfer is treated self-consistently, albeit only approximately, in a 3D spatially-inhomogeneous and time-dependent manner. This is a new addition to the SLH-P^{3}M code. The specific implementation of the radiative transfer is described in the Appendix.

Hereafter I adopt the following values of cosmological parameters in the CDM+Λ model:

$$\Omega_0 = 0.3, \quad \Omega_\Lambda = 0.7, \quad h = 0.7, \quad \Omega_b = 0.04.$$

This cosmological model is in a reasonable agreement with all available observational data and has a desirable feature that all the parameters are specified to only one decimal place. For the simulations presented in this paper the particular choice of cosmological parameters matters rather little as all reasonable models look similar at high redshift.

There are two free parameters that control the behavior of the simulations. The first parameter is the efficiency of star formation ϵ_*, that appears in the Schmidt law:

$$\frac{d\rho_*}{dt} = \epsilon_* \frac{\rho_g}{t_*}, \tag{1}$$

where ρ_* and ρ_g are the stellar and the gas density respectively, and t_* is the maximum of the dynamical and cooling time. Equation (1) is only applied in resolution elements that are determined to be beyond the resolution limit of the simulation. In all resolved resolution elements star formation is not allowed.

The second important parameter measures the amount of ionizing radiation each "star" emits:

$$\frac{du_{UV}}{dt} = \epsilon_{UV} c^2 \frac{d\rho_*}{dt}, \tag{2}$$

where u_{UV} is the production of the energy density in ionizing radiation per unit time. This parameter apparently depends on the initial mass function (IMF). For the Salpeter IMF the "escape fraction" of the ionizing photons, i.e. the fraction of ionizing photons that escape from the immediate vicinity of a star, is roughly

$$f_{\rm esc} \approx 1.4 \times 10^4 \epsilon_{UV} \tag{3}$$

(Madau et al. 1999). Thus, an efficiency of 4×10^{-5} corresponds to roughly 60% escape fraction.

However, what matters most for the evolution of the IGM is only the total number of ionizing photons emitted, i.e. the product of ϵ_* and ϵ_{UV}. Therefore, I fix $\epsilon_* = 0.05$ as my fiducial value, which is similar to the values found in the Milky Way. The latter by itself is not a justification for this particular choice, since there may be very little in common between the star formation in the Milky Way and at high redshift. However, this choice leads to about 3 to 5 percent of all baryons being converted to stars by $z = 4$, which translates to about 10 to 20 percent of all stars being formed by $z = 4$, a reasonable estimate given the current data on the star formation history of the universe (Madau 1999, Steidel et al. 1999, Renzini 1999). In addition, the star formation rate at $z = 4$ in the simulation turns out to be similar to the observational value (see Fig. 2 below).

It is important to emphasize here that an assumed value of the radiation efficiency, or for that matter of the escape fraction, is resolution dependent, since it measures the fraction of energy that escapes from simulated stars into the simulated gas. Since any simulation has a finite resolution, this means that the escape fraction in a simulation measures the fraction of radiation that "escapes from the stellar surface to the resolution scale of the simulation", and mathematically precise meaning of this phrase cannot be formulated, as it may depend on the behavior of a numerical scheme at its resolution limit. It may even be a bad approximation to assume ϵ_{UV} to be constant during the simulation, since simulations presented in this paper have a fixed resolution in comoving coordinates, which means that in physical units the resolution length increases with time. However, since the process of reionization takes a relatively short period of time, this effect is likely to be insignificant. At any rate, as I argue below, the simulations presented in this paper are only reliable on a semi-qualitative level (within a factor of two or so), and cannot be relied upon to provide an accurate quantitative answer.

Table 1 lists all the simulations presented in this paper. The last column in the table is the final redshift at which a given simulation is stopped. The first four simulations whose labels begin with N64 contain 64^3 dark matter particles, 64^3 baryonic mesh, and a number of stellar particles which form during the simulation. These runs have a dynamic range of 1300. All 64^3 runs have a computational box size of $2h^{-1}$ comoving megaparsecs.

The 64^3 runs serve as tests and designed to investigate the parameter dependence and the convergence of the simulations. They are not used for producing scientific results. The first three

Table 1. Numerical Parameters

Run	N	Box size	Baryonic mass res.	Spatial res.	ϵ_*	ϵ_{UV}	z_f
N64_L2_A	64^3	$2h^{-1}$ Mpc	$10^{5.7}\,M_\odot$	$1.5h^{-1}$ kpc	0.05	4×10^{-5}	4
N64_L2_B	64^3	$2h^{-1}$ Mpc	$10^{5.7}\,M_\odot$	$1.5h^{-1}$ kpc	0.05	1.2×10^{-4}	4
N64_L2_C	64^3	$2h^{-1}$ Mpc	$10^{5.7}\,M_\odot$	$1.5h^{-1}$ kpc	0.05	1.3×10^{-5}	4
N64_L2_D	64^3	$2h^{-1}$ Mpc	$10^{5.7}\,M_\odot$	$1.5h^{-1}$ kpc	0.10	4×10^{-5}	4
N128_L4_A	128^3	$4h^{-1}$ Mpc	$10^{5.7}\,M_\odot$	$1.0h^{-1}$ kpc	0.05	4×10^{-5}	4
N128_L8_A	128^3	$8h^{-1}$ Mpc	$10^{6.6}\,M_\odot$	$2.0h^{-1}$ kpc	0.05	4×10^{-5}	6.5
N128_L2_A	128^3	$2h^{-1}$ Mpc	$10^{4.8}\,M_\odot$	$0.5h^{-1}$ kpc	0.025	4×10^{-5}	6.5

of them differ only by the value of the radiation efficiency, and the last run, N64_L2_D, has a twice higher star formation rate. One may note that the escape fraction in run N64_L2_B is formally 170% if equation (3) is used. This should not be considered unphysical since this run is only used as a test, and besides that a 70% increase in the fraction of high mass stars will compensate for that. Anyway, as I show below, this value of the radiation efficiency does not fit the observations.

The last three simulations listed in Table 1 whose labels begin with N128 incorporate 128^3 dark matter particles, the same number of baryonic cells, and about twice more stellar particles, which keep forming during the simulation. The dynamic range of these three simulations is 4000, and other parameters are listed in the table. Run N128_L4_A is the production run of this paper, and two other 128^3 runs are also used to investigate numerical convergence, and therefore are not continued until $z = 4$.

Figure 1 illustrates in a graphical form the mass and spatial scales simulated in this paper. The production run N128_L4_A has just enough mass resolution to resolve the filtering scale, the minimum scale on which baryons cluster in linear theory.[2] This is clearly only marginally enough, as one would prefer to resolve the filtering scale by a factor of 10 or more to have a safe margin. On the other hand the production run N128_L4_A does not miss a large amount of small scale power. Thus, even from this consideration we should expect that the simulations presented in this paper are not quantitatively accurate, but should be sufficient to give a general qualitative picture, since they include all the relevant physics and account for most of the initial power.

Figure 2 now demonstrates the level of convergence of the simulations presented in this paper and the method of choosing the parameter ϵ_{UV}. Three thin lines: dotted, short-dashed, and

[2]As has been shown in Gnedin & Hui (1998) at $z = 10$ the mass corresponding to this scale is about 11 times larger than the Jeans mass at this redshift.

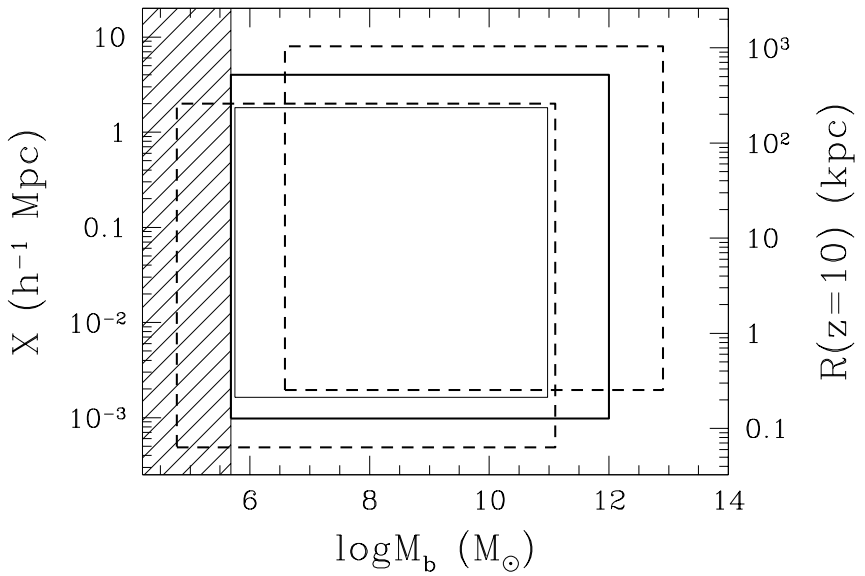

Fig. 1.— Mass and spatial resolution of the simulations presented in Table 1. Left y axis labels comoving scales and right y-axis labels physical scales at $z = 10$. The thin solid box shows all four 64^3 runs (slightly offset for clarity), the bold solid box shows run N128_L4_A, and two bold dashed boxes show runs N128_L2_A and N128_L8_A. The shaded regions marks the mass scales on which baryons do not cluster in linear theory due to finite pressure (the so-called filtering scale; Gnedin & Hui 1998).

long-dashed show runs N64_L2_C, N64_L2_A, and N64_L2_B respectively, which only differ by the value of ϵ_{UV}. One can see that at $z = 4$ the value of the ionizing intensity J_{21} is 0.06, 0.4, and 2.5, respectively, changing by a factor of about 6.5 whenever the radiation efficiency changes by a factor of 3. Since the observational values for J_{21} at $z = 4$ range between 0.2 and 0.4 (Lu et al. 1996), I choose $\epsilon_{UV} = 4 \times 10^{-5}$ (used in run N64_L2_A) as a reasonable value for this parameter. Since my simulations cannot be used for accurate quantitative predictions, such a rough agreement with the observations is sufficient.

A similar exercise with 128^3 simulations is however not feasible, as it would require an excessive amount of computer time to run three 128^3 simulations up to $z = 4$ with three different values of ϵ_{UV}. However, as can be seen from the top panel of Fig. 2, the production run simulation N128_L4_A with assumed value of $\epsilon_{UV} = 4 \times 10^{-5}$ also gives $J_{21} = 0.4$ at $z = 4$, and thus a guess based on 64^3 simulations turns out to be a good one, whereas a small run N64_L2_D, which has twice higher star formation rate than N64_L2_A, agrees better with the production run before and during the epoch of overlap (indicated by a sharp rise in J_{21}).

The two other 128^3 runs were normalized so as to agree with the production run N128_L4_A at $z \sim 10$ in the star formation rate and the ionizing intensity J_{21} (the small box run N128_L2_A actually has somewhat lower star formation rate). Nevertheless, they disagree with it at the epoch of overlap, and this difference indicates the lack of numerical convergence.

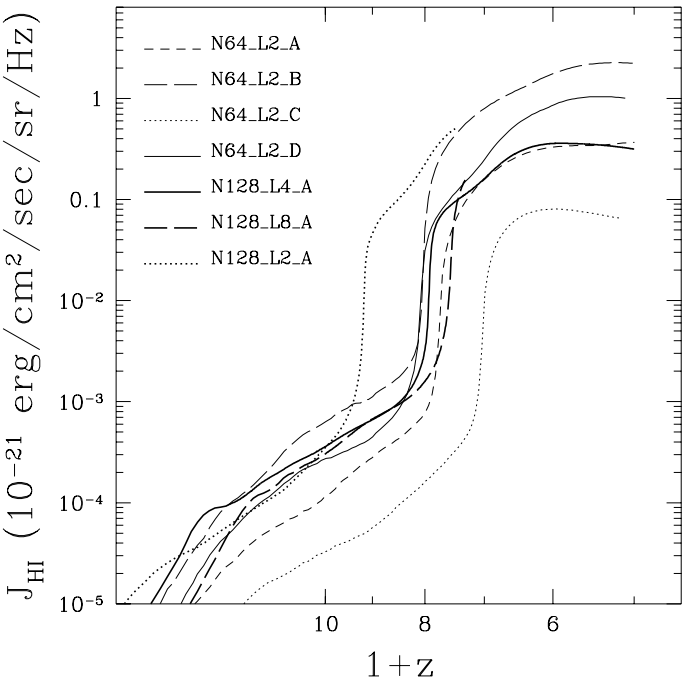

Fig. 2a

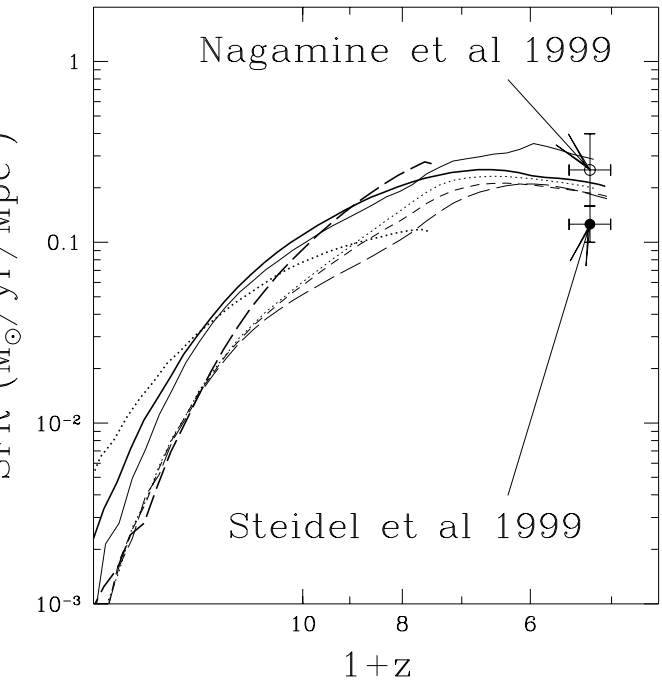

Fig. 2b

Fig. 2.— The evolution of the spatially averaged ionizing intensity J_{21} (measured in conventional units of 10^{-21} erg/ cm^2/ sec/ Hz/ rad) (*a*) and the star formation rate (*b*) as a function of redshift for seven runs listed in Table 1: N64_L2_A (*thin short-dashed line*), N64_L2_B (*thin long-dashed line*), N64_L2_C (*thin dotted line*), N64_L2_D (*thin solid line*), N128_L4_A (*bold solid line*), N128_L8_A (*bold long-dashed line*), N128_L2_A (*bold dotted line*). The data points for the star formation rate at $z \approx 4.2$ are from Steidel et al. (1999; *filled circle*) and the same data corrected for dust extinction by a different method by Nagamine, Cen, & Ostriker (1999; *empty circle*).

Let me first focus on an $8h^{-1}$ Mpc run N128_L8_A. It agrees quite well with the production ($4h^{-1}$ Mpc) run before the overlap, but predicts a somewhat lower redshift of overlap. It however agrees better with the smaller N64_L2_A run, because its spatial resolution is closer to the small run than to the production run. Thus, the redshift of overlap is delayed if the spatial resolution of a simulation is not sufficient. Is the resolution of the production run sufficient in this respect? We can use comparison between the production run N128_L4_A and a $2h^{-1}$ Mpc run N128_L2_A to investigate the effect of spatial resolution. It is clear that the production run still lacks resolution to resolve the very first star formation at $z \sim 20$. This is not surprising, since it barely resolves the baryon filtering scale before the reheating, and thus misses about 7% of the total small scale power. At $z \sim 10$ the two runs more-or-less agree, but the smaller run has a much earlier time of overlap. Also, it gives J_{21} at $z = 6.5$ of about 0.6, already higher than the final value of the production run of about 0.4. Thus, run N128_L2_A can be considered as giving the upper limit for the redshift of overlap for this model (which is $z = 8.5$), and this upper limit is likely to be too high, because run N128_L2_A is going to have a too high a value for J_{21} at lower redshifts.

This is of course not surprising, since run N128_L2_A has the same box size as N64_L2_A, but much higher spatial and mass resolution, so it is going to have more star formation than a smaller run, and thus an earlier reionization. In comparison to the production run N128_L4_A, it misses important large scale waves which contribute toward absorption. It also suffers from the cosmic variance problem, since at $z = 4$ the rms density fluctuation in a $2h^{-1}$ Mpc box is 0.41, whereas for a $4h^{-1}$ Mpc it is 0.27, and for a $8h^{-1}$ Mpc it is 0.16.

The comparison above does not however constitute the complete test of numerical convergence, as three different resolutions are required for such a test. Thus, a 256^3 run is needed to assign a meaningful value for the numerical error of my simulations, but such a run is currently beyond the available means. I must therefore acknowledge that the simulations presented in this paper can only be considered reliable on a semi-qualitative basis, within a factor of two or so. One can use the difference between the bold solid line and thin solid and dashed lines as an estimate of the numerical error, but without a larger simulation it is impossible to assign any meaningful confidence level to this error.

3. Results

3.1. Reionization at a glance

We can now look at a general description of the reionization process, as reflected in the simulations. In order to visualize the process, I show in Figure 3 a thin ($15h^{-1}$ comoving kiloparsecs deep) slice through the computational box taken at a random place. The three panels show in a logarithmic stretch the neutral hydrogen (upper left panel), the gas density (lower left panel), and the gas temperature (lower right panel). The upper right panel shows the evolution of the ionizing intensity (logarithmically scaled) versus the redshift (the same as the solid bold line in Fig. 2a). The scale bars next to respective panels show the correspondence between the color and the numeric scale. The epoch as measured by the scale factor $a \equiv 1/(1 + z)$ is labeled at the center of each plot.[3]

I can now describe the general features of reionization. The reionization starts (Fig. 3a) with ionization fronts propagating from proto-galaxies located in high density regions into the voids, leaving the high density outskirts of an object still neutral, because at high density the recombination time is very short, and there are not enough photons to ionize the high density regions. As an ionization front moves on with the H II region forming behind it, it leaves behind high density regions which require many more photons to get ionized than is available at that moment (Fig. 3a-d). This stage of the reionization process can be called "pre-overlap", and it extends over a considerable range of redshifts $\Delta z \sim 5$ around $z \sim 10$. During this time the high

[3]An MPEG movie of the simulation is available at `http://casa.colorado.edu/~gnedin/GALLERY/rei_p.html`.

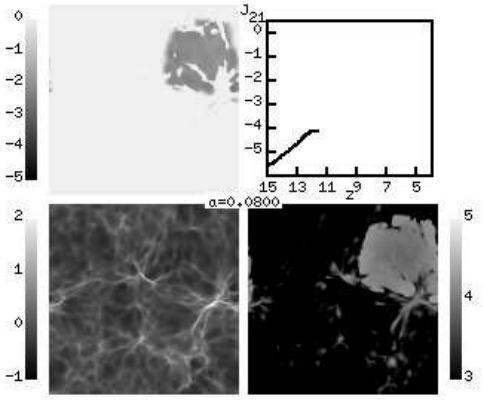

Fig. 3a

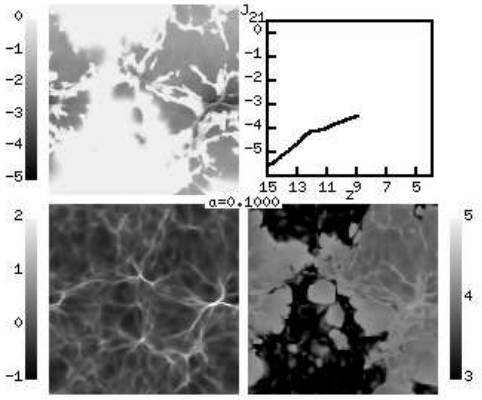

Fig. 3b

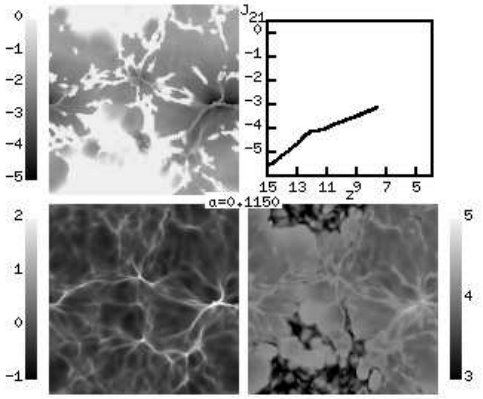

Fig. 3c

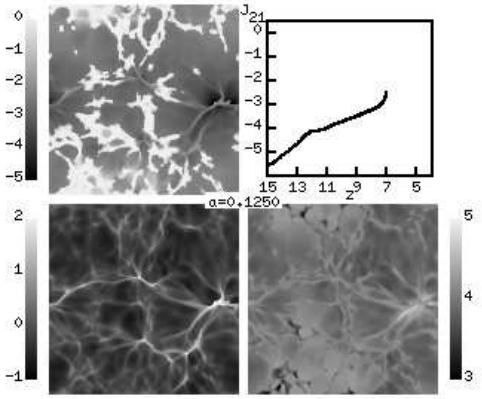

Fig. 3d

Fig. 3.— A thin slice through the simulation volume at eight different epochs: (*a*) $z = 11.5$, (*b*) $z = 9$, (*c*) $z = 7.7$, (*d*) $z = 7$, (*e*) $z = 6.7$, (*f*) $z = 6.1$, (*g*) $z = 5.7$, (*h*) $z = 4.9$. Shown are logarithm of neutral hydrogen (*upper-left*), logarithm of gas density (*lower-left*), logarithm of gas temperature (*lower-right*), and $\lg(J_{21})$ as a function of redshift (*upper-right*).

density regions around the source are slowly becoming ionized, whereas high density regions far from the source remain neutral. The ionizing intensity at this time remains low and is slowly increasing with time. However, since the radiation field is highly inhomogeneous at this time, the ionizing intensity J_{21}, which is by definition a space-average, has only a formal meaning.

By $z \approx 7$ the H II regions start to overlap (Fig. 3d), with the result that the number of sources shining at an average place in the universe increases, and the ionizing intensity starts to rise rapidly. The process of reionization enters its second stage, the "overlap", which is quite rapid (Fig. 3d-e). As the ionizing intensity is rapidly increasing, the last remains of the neutral low density IGM are quickly eliminated, the mean free path increases by some two orders of magnitude (as explained below) over a Hubble time or so, and voids become highly ionized (neutral fraction of the order of 10^{-5}). The high density regions at this moment are still neutral, as the number of ionizing photons available is not sufficient to ionize them.

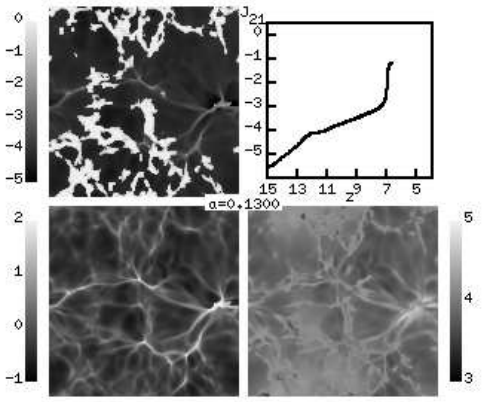

Fig. 3e

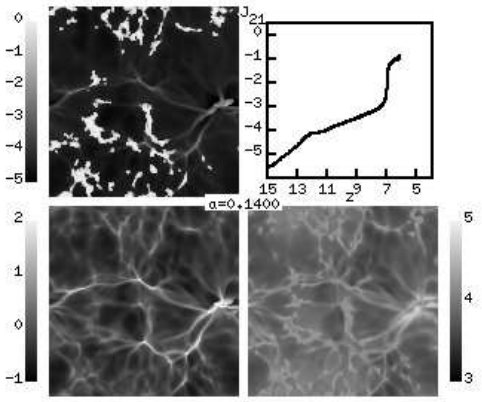

Fig. 3f

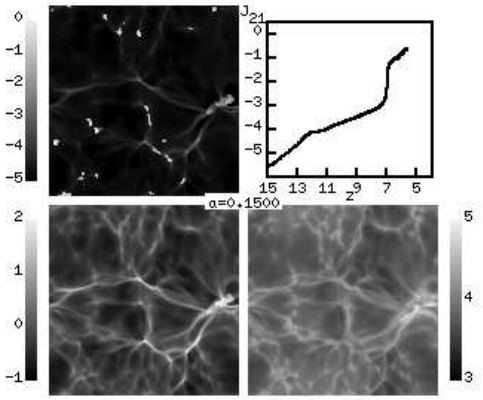

Fig. 3g

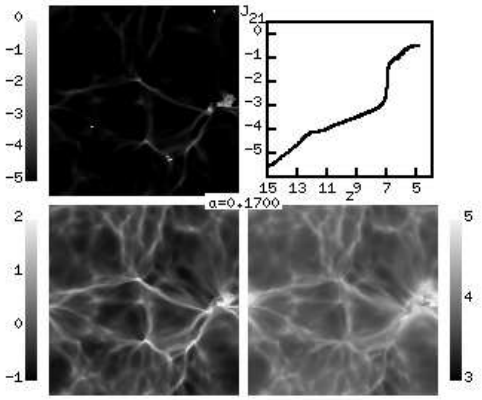

Fig. 3h

After the overlap is complete, the universe is left with highly ionized both low density regions and some of the high density ones (which happened to lie close to the source, where the local value of the ionizing intensity is higher than the spatial average). High density regions far from any source remain neutral. This stage can be called "post-overlap", and is well described by the semi-analytic model of Miralda-Escudé et al. (1999). As time goes on, and more and more ionized photons are emitted, the high density regions are gradually being eaten away (Fig. 3e-g), and the spatially averaged ionizing intensity J_{21} continues to rise slowly until it flattens out at $z \sim 5$ (Fig. 3g). The latter effect is however likely to be an artifact of a finite simulation box, as by this time the mean free path exceeds the box size by about a factor of 10 (as is shown below), and the simulation runs out of high enough density peaks, which would be present in the real universe (or a larger box simulation), and fails to reproduce the reality even on a semi-qualitative level.

Therefore, the recent argument of whether reionization is "fast" or "slow" (see, for example, Miralda-Escudé et al. 1999 versus Gnedin & Ostriker 1997 or Madau et al. 1999) is, in large part, a question of terminology. If one considers the whole process of reionization, which consists of "pre-overlap", "overlap", and "post-overlap", one inevitably concludes that reionization is slow (i.e. taking place over a Hubble time or longer). However, if one looks only at the process of

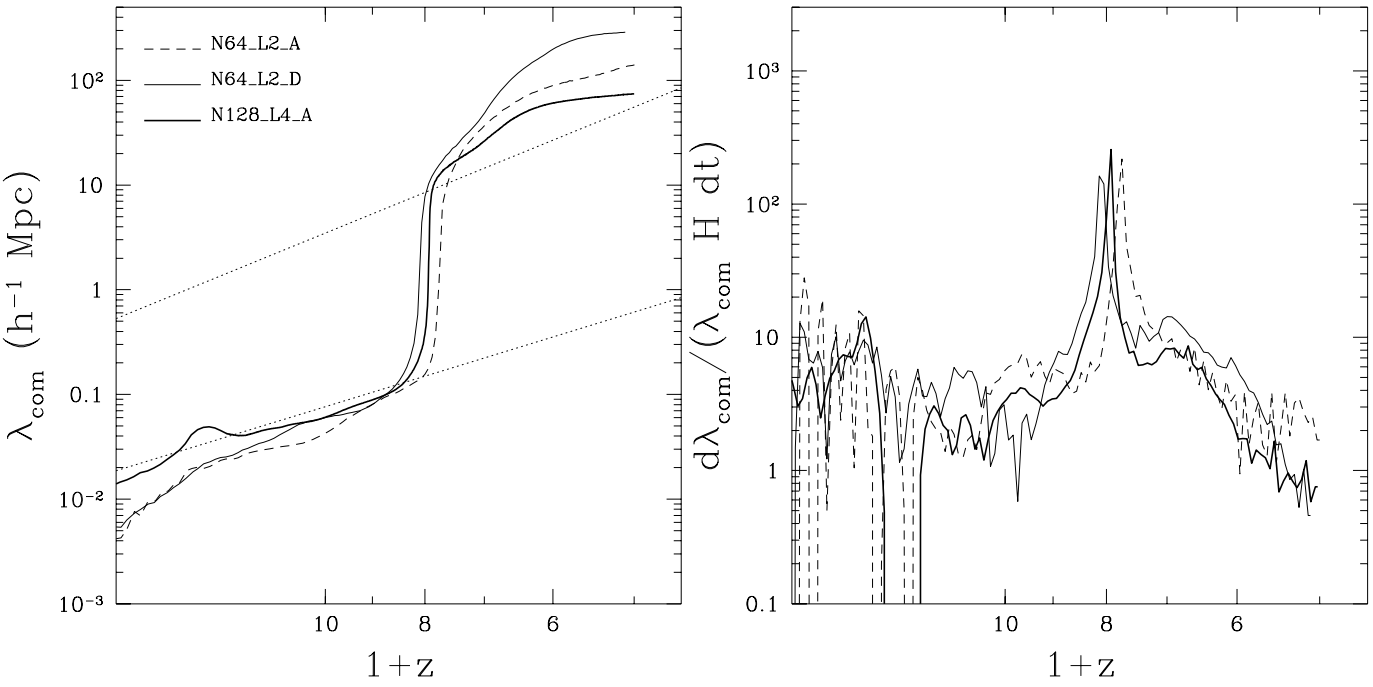

Fig. 4a Fig. 4b

Fig. 4.— The evolution of the mean free path (a) and its rate of change (b) for the production run N128_L4_A (*bold solid line*) and two smaller runs N64_L2_D (*thin solid line*) and N64_L2_A (*thin short-dashed line*). Two thin dotted lines are explained in the text.

"overlap" and labels it reionization, then one concludes that reionization is fast.

3.2. Reionization in more detail

Let me now focus on details of the process that is overviewed in the previous subsection. Figure 4 shows the evolution of the mean free path and its time derivative for the production run N128_L4_A and two smaller runs (the latter two to demonstrate the level of numerical convergence). The epoch of overlap is clearly distinguished by a sharp rise in the mean free path. The rate of this rise, some 100 times faster than the expansion of the universe, shows no trend of diminishing with the increase of the box size (the bold line versus thin ones), and thus is unlikely to be an artifact of a finite simulation box.

Can this behavior be understood in simple terms? Let me consider a simple model for the universe: each object has a $1/r^2$ density profile around it until those profiles overlap with neighboring objects. A source, placed at the center of an object produces the photoionization rate that falls off as $1/r^2$ in the optically thin regime,

$$\Gamma(r) = \Gamma_0 \frac{r_0^2}{r^2}.$$

Given the density profile with a core radius r_0,

$$n(r) = n_0 \frac{r_0^2}{r_0^2 + r^2},$$

the neutral hydrogen fraction at $r \gg r_0$ is constant,

$$x(r) = x_0 = \frac{Rn(r)}{\Gamma(r)} = \frac{Rn_0}{\Gamma_0}.$$

The mean free path then is given by the following expression,

$$\lambda = \frac{c}{\bar{k}},$$

where $\bar{k}$ is the photoionization rate averaged value of the absorption coefficient (I omit here frequency dependence for simplicity):

$$\bar{k} = \frac{\langle k\Gamma \rangle}{\langle \Gamma \rangle} \tag{4}$$

(see Gnedin & Ostriker 1997, eqn. [A6]), and $\langle \Gamma \rangle$ is the photoionization rate spatially averaged over a sphere of radius R,

$$\langle \Gamma \rangle = \frac{3}{R^3} \int_0^R \Gamma(r) r^2 dr = 3\Gamma_0 \frac{r_0^2}{R^2}.$$

Using the expression for the density, for the mean free path I obtain:

$$\lambda = \frac{2}{\pi \sigma x_0 n_0} \frac{R}{r_0}, \tag{5}$$

where σ is the ionization cross section. If, instead of $\bar{k}$, I used a spatially averaged absorption coefficient $\langle k \rangle$, which would be appropriate if the ionizing background was homogeneous, I would get a different result:

$$\hat{\lambda} \equiv \frac{c}{\langle k \rangle} = \frac{1}{3\sigma x_0 n_0} \frac{R^2}{r_0^2}. \tag{6}$$

Reasonable values for n_0 and r_0 would be the virial density, $n_0 = 200\bar{n}_b = 4 \times 10^{-5}(1+z)^3\,\mathrm{cm}^{-3}$, where $\bar{n}_b$ is the average baryon density, and the virial radius, $r_0 = 31(1+z)^{-1}\,\mathrm{kpc}$ for a 10^9 solar mass object. Obviously, R has to be about the size of the H II region around the source, which from Fig. 3 can be roughly approximated as $R = 10(1+z)^{-2}\,\mathrm{Mpc}$ at $z \sim 9$. I also assume $x_0 \sim 10^{-2}$, which is justified below.

The two dotted lines in Fig. 4a now show λ (the lower line) and $\hat{\lambda}$ (the upper line). Before the overlap, the radiation field is highly nonuniform, and the mean free path is short, as the majority of photons are absorbed within one H II region. The lower dotted line (λ from eqn. [5]) provides in this regime a remarkably good fit to the simulation results despite the highly simplistic nature of the above derivation. Fig. 3b-d illustrate this effect. The H II region located in the middle right of the slice contains the high density neutral filaments, which are slowly being eaten away by the ionization front, which moves much slower across high density filaments than over the low density

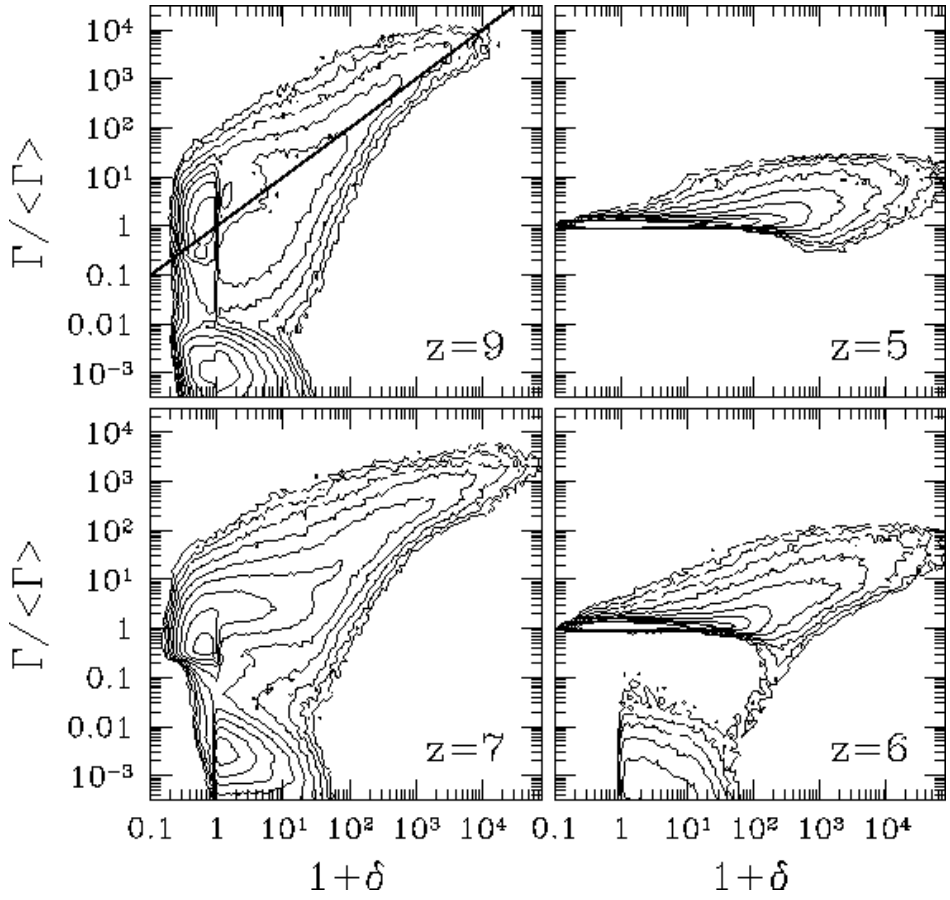

Fig. 5.— The joint mass-weighted distribution of the gas density and the local photoionization rate Γ in the production run N128_L4_A at four different redshifts. The straight bold line in the upper left panel marks a $\Gamma = \langle\Gamma\rangle(1+\delta)$ law.

void. The absorption in these high density regions shortens the mean free path compared to the size of the H II region in the low density medium. This is consistent with the simplistic picture above which assumes a spherically symmetric density distribution, which in reality is an average of the high density filaments (with the density run shallower than $1/r^2$) and the low density voids (with the density run steeper than $1/r^2$).

After the overlap, the ionizing radiation is more (albeit non perfectly) homogeneous, and the mean free path increases by almost two orders of magnitude and becomes comparable to the estimate $\hat{\lambda}$. Thus, the effect of the overlap being "fast" is entirely due to the change in the relative distribution of the ionizing radiation and the gas density, and thus, in essence, is a topological effect.

This is further illustrated by Figure 5, which shows on four panels the joint mass-weighted distribution of the gas density and the photoionization rate at four different redshift. At $z = 9$ the photoionization rate in the regions that are indeed photoionized (only those count towards the

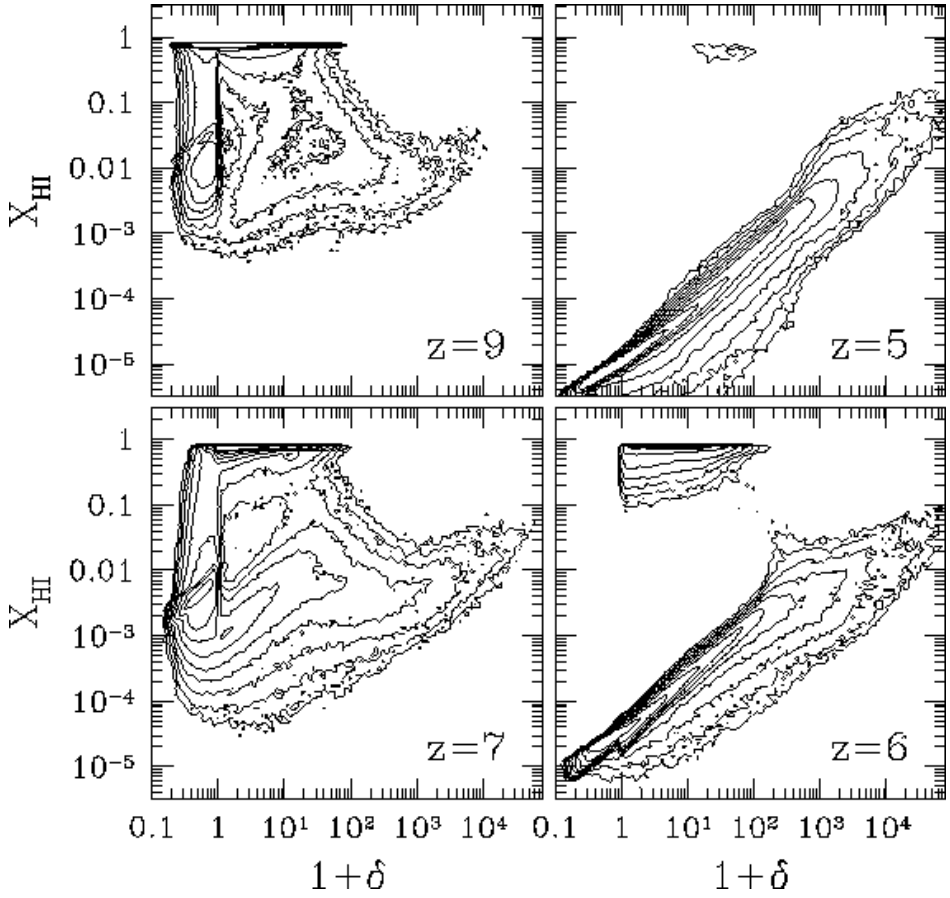

Fig. 6.— The joint mass-weighted distribution of the gas density and neutral hydrogen fraction in the production run N128_L4_A at four different redshifts. The excessive sharpness of the features at $\delta \approx 0$ is the defect of the radiative transfer approximation, in reality there should be a more gradual transition from the underdense to the overdense regions.

mean free path, since no photons are absorbed where there are no photons) is proportional to the gas density, with a large scatter reflecting the fact that different objects have different luminosities (in terms of the simplistic model above, Γ_0 is different in different objects). As the H II regions start to overlap ($z = 7$), regions with more or less uniform photoionization rate starts to appear at low densities. Finally, in the post-overlap stage, a large fraction of the total volume has a quite uniform photoionization rate, which still varies significantly in the high density regions ($z = 6$ and $z = 5$). I will address the question of the inhomogeneity in the photoionization rate at low redshifts below.

Figure 6 illustrates the process of reionization from yet another point of view. It presents the mass-weighted joint distribution of the gas density and neutral hydrogen fraction shown at the same four redshifts $z = 9$, 7, 6, and 5. In the pre-overlap stage the neutral hydrogen fraction is roughly independent of density in the photoionized regions, again with a large spread due to

different luminosities of different objects and with the median value of about $x_{\rm H\,I} \sim 10^{-2}$, which is a justification for adopting this value in the simplistic estimate presented above. As time progresses and the H I regions start to overlap, a regime $x_{\rm H\,I} \propto (1+\delta)$ starts to develop in the lower density gas, which is just a reflection of the fact that the photoionization rate is becoming independent of the density in those regions, whereas the high density gas still maintains the neutral fraction that is roughly independent of the density. After the overlap, most of the gas sits at the ionization equilibrium in a tight relation with the density, which becomes somewhat more spread in the highest density regions which correspond to the inner parts of individual objects. However, a considerable fraction of the high density gas remains in the neutral state, in accordance with the semi-analytical picture of Miralda-Escudé et al. (1999). As the production of the ionizing photons continues, this gas is slowly being ionized, mostly the lower density regions first.

One of the outcomes of the main assumption of this paper - that the sources of ionization are proto-galaxies rather than quasars - is that each halo harbors a source (except, perhaps, very low mass ones, that are not well resolved in my simulation), and therefore the halos are also ionized by local sources. The resolution of my simulation is however insufficient to resolve the interstellar medium, and proto-galaxies will contain giant molecular clouds which will be responsible for the damped Lyman-alpha absorption.

3.3. What does it?

The understanding of the process of reionization would not be complete without identifying which sources are indeed responsible for the bulk of ionizing photons. To this end, I have identified all bound objects in the production run N128_L4_A at $z = 7$ (the epoch of overlap) using Bertchinger & Gelb's (1991) DENMAX algorithm. Figure 7a shows the star formation rate versus the stellar mass for all objects. As can be seen, there exists a good correlation between the star formation rate and the mass of the object,

$$\dot{M}_* \approx \frac{M_*}{1.6 \times 10^8\,{\rm M}_\odot} \frac{{\rm M}_\odot}{\rm yr},$$

reflecting the fact that all objects found in the simulation form stars with the rapidly increasing rate (Fig. 2).

However, not all objects contribute equally toward the reionization of the low density IGM, but only those whose ionization fronts are able to leave their own halo. While the detailed analysis of the evolution of each individual ionization front is beyond the scope of this paper, a simple picture can be introduced that captures the main qualitative features of the simulation.

If each individual object can be approximated as a homogeneous sphere with the radius equal the virial radius of the object, and the mean overdensity of 200, than an object will contribute towards the reionization of the low density IGM if the virial radius r_v is smaller than the Strömgen radius r_s of the ionization front. Figure 7b now shows the the Strömgen radius versus the virial

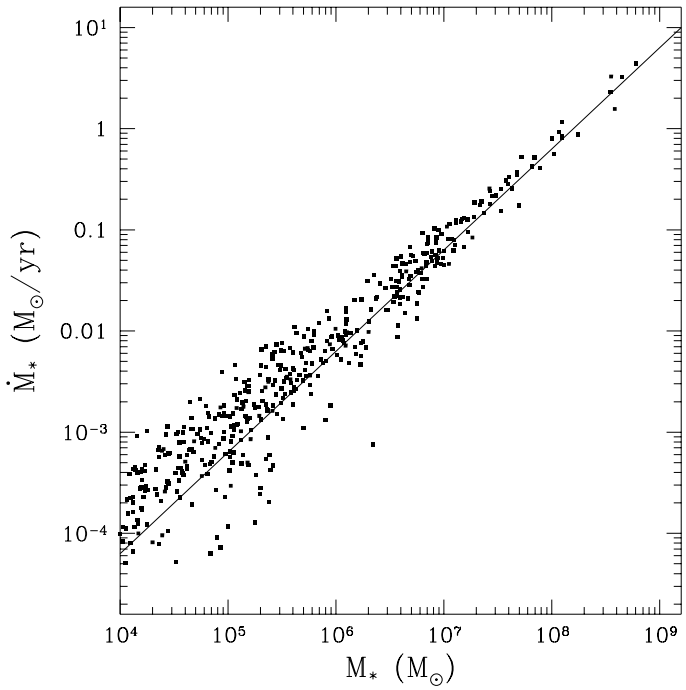

Fig. 7a

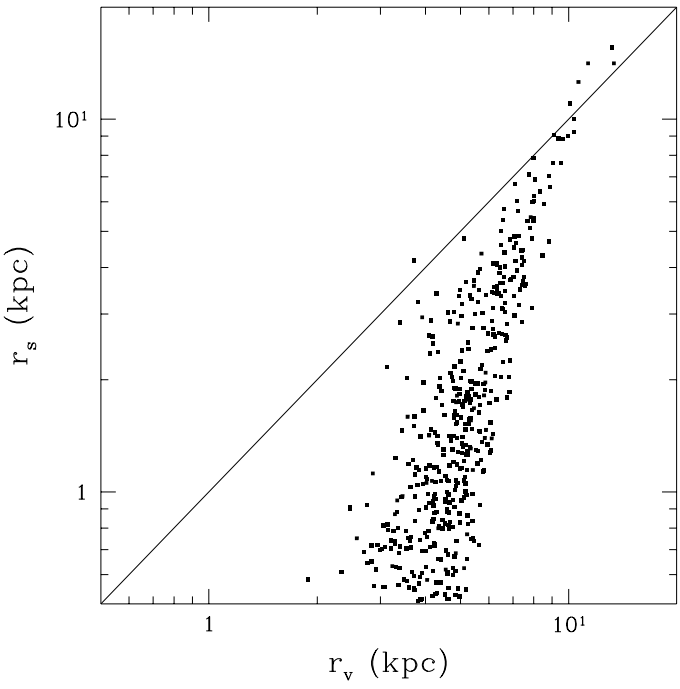

Fig. 7b

Fig. 7.— The star formation rate versus the stellar mass (a) and the virial radius versus the Strömgen radius (b) for all bound objects in the production run N128_L4_A.

radius for all the bound resolved objects in the production run. One can see that the majority of objects have their Strömgen spheres well inside the virial radius, and only a handful of objects contribute to the ionizing flux that is capable of ionizing the low density IGM.

This is not surprising, given what is seen in Fig. 3, and the fact that I have not been able to establish numerical convergence - if only a few objects in my simulation box are responsible for reionization, then the $4h^{-1}$ Mpc box size is too small to be a representative volume of the universe, and a larger box size is needed to achieve the full numerical convergence.

Figure 8 serves to illustrate this further. It shows the stellar mass function and the star formation rate function for the bound objects from my simulation, as well as the Schecter function fit to the stellar mass function (the fit deviates at low masses due to photoionization feedback). In addition, I plot with the dotted line the modified star formation rate function which includes only photons capable of ionizing the low density IGM, defined as

$$\dot{M}_I \frac{dn}{d\dot{M}_I},$$

where

$$\dot{M}_I = \max(0, (r_s/r_v)^3 - 1)\dot{M}_*$$

is proportional to the number of photons that escape into the IGM. As can be seen from Fig. 8, it

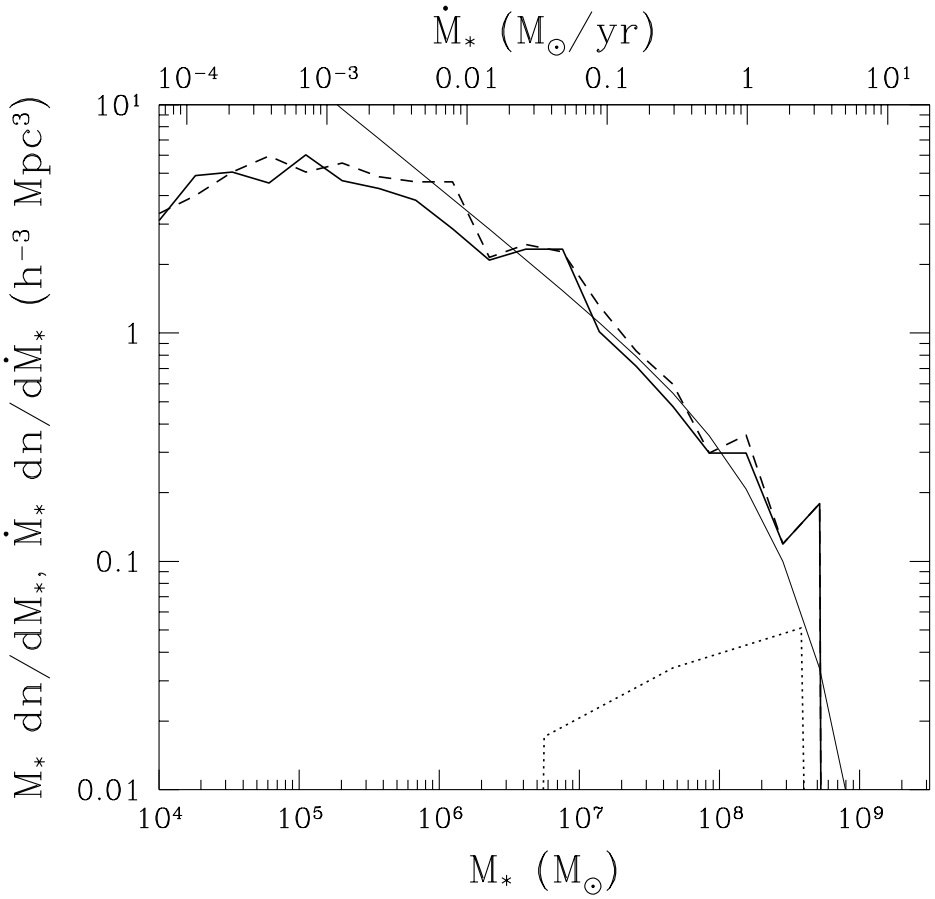

Fig. 8.— The stellar mass function (*solid bold line*) and the star formation rate function (*dashed line*) at $z = 7$. The dotted line shows the modified star formation rate function which includes only photons capable of ionizing the low density IGM. The thin solid line is a Schecter function fit to the stellar mass function with $\alpha = -0.5$ and $M_\star = 3 \times 10^8\,\mathrm{M}_\odot$.

is objects with luminosities in excess of $0.1L_\star$ that are responsible for reionization, with $L_\star$ and more luminous objects making the main contribution.

3.4. Comparison with the semi-analytical models

I have mentioned already that my simulations agree with the qualitative predictions of Miralda-Escudé et al. (1999) quite well in the post-overlap stage. Figure 9 serves to illustrate the agreement even further. Fig. 9a shows a comparison between the mean volume emissivity ϵ and the global recombination rate R, as introduced in equation (2) of Miralda-Escudé et al. (1999). In accordance with their prediction, $R \approx \epsilon$ just after the overlap, and continues to be close to ϵ and somewhat lower when other terms become important.

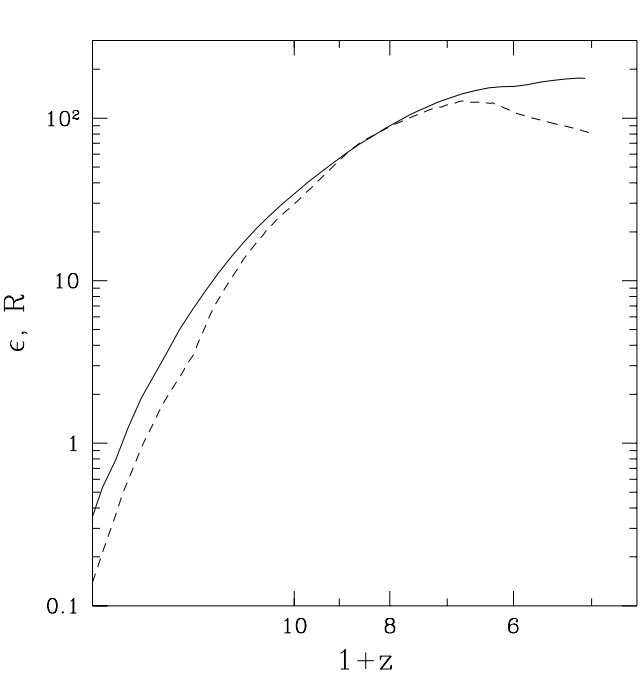

Fig. 9a

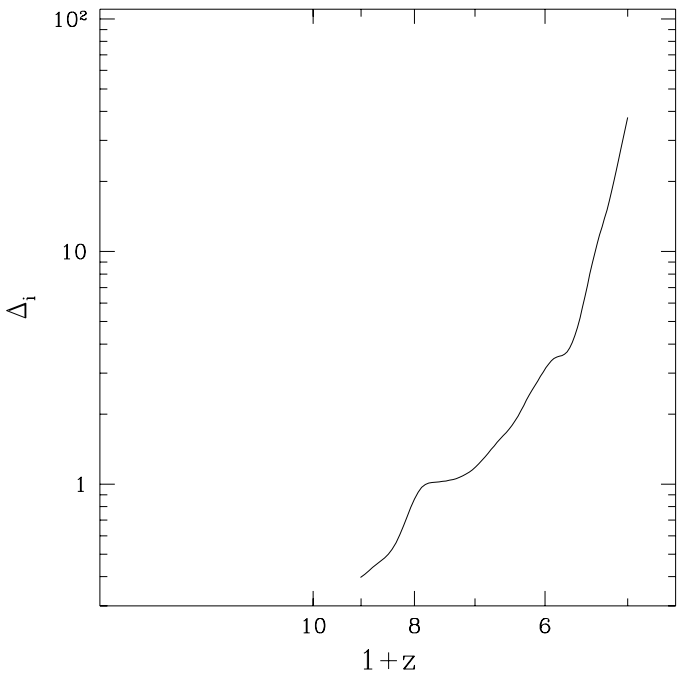

Fig. 9b

Fig. 9.— (*a*) Comparison of the mean volume emissivity ϵ (*solid line*), measured in terms of the number of ionizing photons emitted for each atom in the universe per Hubble time, and the global recombination rate R (*dashed line*), measured in terms of the mean number of recombinations per Hubble time per baryon, as introduced in equation (2) of Miralda-Escudé et al. (1999). (*b*) The time evolution of the quantity Δ_i as defined in equation (1) of Miralda-Escudé et al. (1999).

Fig. 9b shows the evolution of the quantity Δ_i defined in equation (1) of Miralda-Escudé et al. (1999). Physically, Δ_i measures the cosmic density (in units of the mean density) above which the gas is mostly neutral, and below which the gas is mostly ionized. While Fig. 6 demonstrates that at high density the ionization state of the gas is not related to the density in any simple way, but also depends on whether a given fluid element is close or far from the nearest source, it is still possible to define Δ_i after the overlap as the limiting density above which 95% of the neutral gas lies. This quantity is plotted in Fig. 9b. While specific numerical values for Δ_i are significantly lower than those quoted in Miralda-Escudé et al. (1999), because there is also the high density ionized gas responsible for a large number of recombinations, the general trend of Δ_i increasing with time is clearly observed.

Figure 10 shows the evolution of the mean number of ionized photons per baryon $N_{\gamma/b}$ on top and the neutral hydrogen fraction at the bottom for the production run and for the corresponding small run. One can see that by the time of overlap $z = 7$, about 10 ionizing photons were emitted per baryon, which may be considered a contradiction with Miralda-Escudé et al. (1999). I must however reemphasize here that the total number of photons emitted depends on the radiation

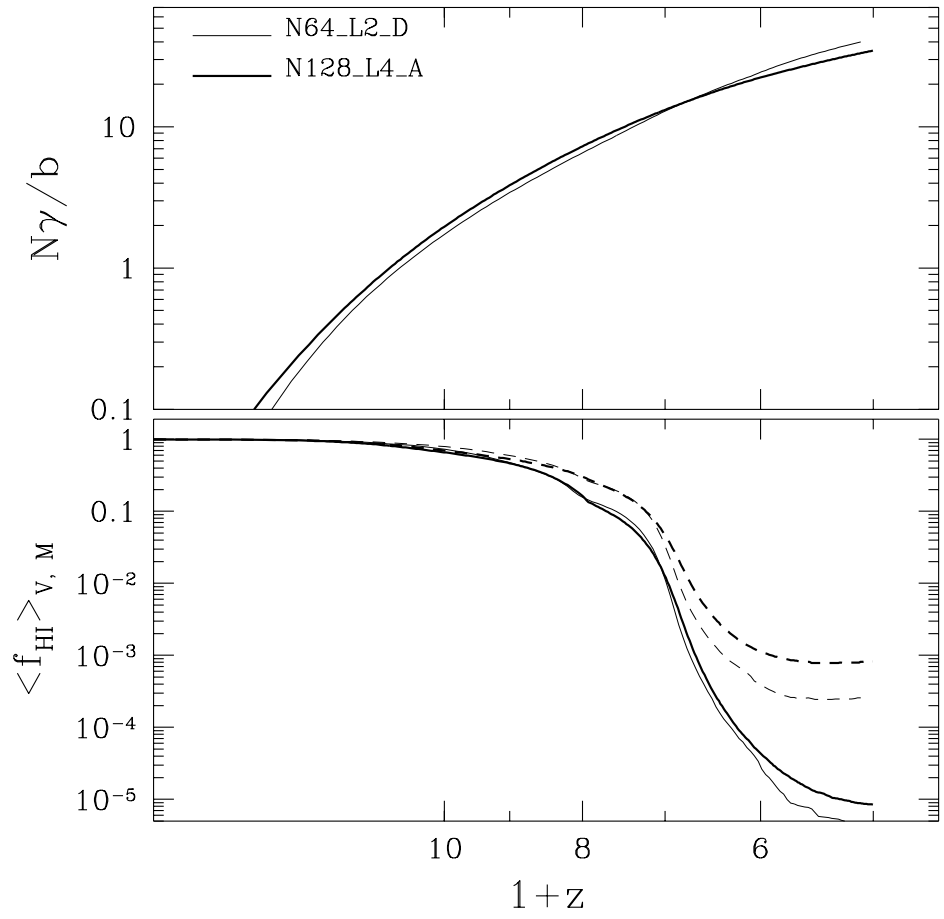

Fig. 10.— *Top panel*: the mean number of ionizing photons emitted per baryon as a function of redshift for the production run N128_L4_A (*bold line*) and a smaller run N64_L2_D (*thin line line*). *Bottom panel*: the mass- (*dotted lines*) and volume-averaged (*solid lines*) neutral hydrogen fractions for the production run N128_L4_A (*bold lines*) and a smaller run N64_L2_D (*thin solid lines*).

efficiency ϵ_{UV}, which in turn is highly resolution dependent. Therefore the quantity $N_{\gamma/b}$ is also resolution dependent and has no physical meaning taken at a face value. The quantity used by Miralda-Escudé et al. (1999) is however the number of ionizing photons per baryon escaped into the IGM from the immediate neighborhood of the ionizing sources,[4] and is therefore considerably smaller than the number directly taken from the simulation. The latter quantity cannot be much less than 1, and since it has to be considerably smaller than 10, the best estimate one can make is that it is of order of a few, again in agreement with Miralda-Escudé et al. (1999).

A large number of workers has studied reionization within the framework of the semi-analytical modeling using the clumping factor approach (see the references in the Introduction). Their

[4]This quantity is physically intuitive, but is difficult to define in a mathematically rigorous fashion.

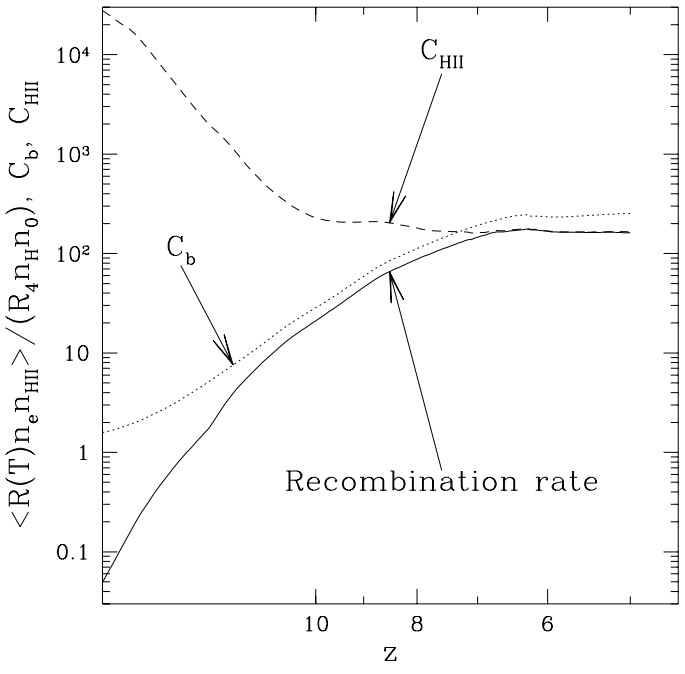

Fig. 11a

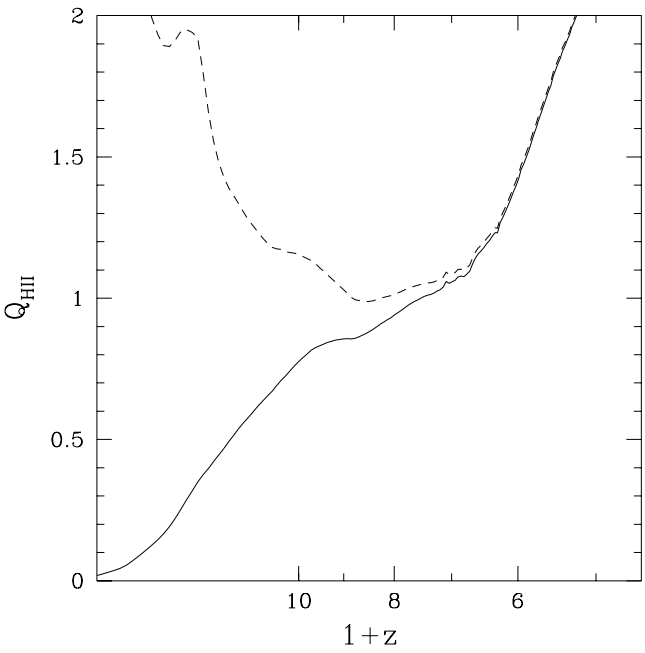

Fig. 11b

Fig. 11.— (*a*) The mean global recombination rate in units of the recombination rate of the fully ionized homogeneous universe (*solid line*), the density clumping factor (*dotted line*), and the ionized hydrogen clumping factor (*dashed line*) as a function of redshift. (*b*) The porosity of the IGM as a function of redshift for the approximation based on equation (23) of Madau et al. (1999) (*dashed line*) and the full calculation based on their equation (21) (*solid line*).

main conclusion is that reionization is fast (because the recombination time is short), and that many ionizing photons per baryon are required to ionize the universe (again because of the same reason). It is therefore useful to compare this approach to the simulations and to Miralda-Escudé et al. (1999). The latter comparison is especially instructive since there appears to be a complete disagreement between the two semi-analytical approaches: if Miralda-Escudé et al. (1999) claim that reionization has to be gradual, the clumping factor approach predicts a fast reionization; if Miralda-Escudé et al. (1999) show that just a few (perhaps as low as one) ionizing photons per baryon are needed for the overlap, the clumping factor approach requires that this number be large.

I now show in Figure 11a the evolution of the volume averaged recombination rate $\langle R(T) n_e n_{\rm H\,II} \rangle$ in units of the recombination rate for the fully ionized homogeneous universe with $T = 10^4\,{\rm K}$ ($R_4 \equiv R(10^4\,{\rm K})$, $n_{\rm H}$ is the total number density of atomic and ionized hydrogen, and $n_0 = n_{\rm H} + n_{\rm He}$, thus assuming that helium is only singly ionized). Also shown the density clumping factor

$$C_b \equiv \frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2},$$

and the ionized hydrogen clumping factor

$$C_{\rm H\,II} \equiv \frac{\langle R(T) n_e n_{\rm H\,II} \rangle}{R_4 \langle n_e \rangle \langle n_{\rm H\,II} \rangle},$$

which is simply related to the volume averaged recombination rate,

$$\langle R(T) n_e n_{\rm H\,II} \rangle = R_4 \bar{n}_e \bar{n}_{\rm H\,II} C_{\rm H\,II}, \tag{7}$$

where bar means volume average.

There are few things to note. First, the recombination rate increases fast with time, and at the epoch of overlap already approaches 100 times that of the homogeneous universe. The density clumping factor is only slightly higher than the recombination rate, and the ionized hydrogen clumping factor begins its evolution with very high values, and approaches the recombination rate at later times, as can be expected from equation (7). This behavior is quite different from what was predicted in Gnedin & Ostriker (1997), and the reason for the disagreement is quite clear: since the ionization fronts originate in the highest density regions, the first H II regions are extremely overdense, which results in very large values for the clumping factor. However, since these first H II regions occupy only a small fraction of the volume, the recombination rate is not large.

Thus, the simulations confirm the predictions of the clumping factor approach: the clumping factor is indeed large at the time of overlap. To do a more detailed comparison, I picked up a Madau et al. (1999) paper as an example of semi-analytical modeling using this type of approach. There have been recently a series of detailed papers on using this method (Haiman & Loeb 1997, 1998; Valageas & Silk 1999; Chiu & Ostriker 1999), but Madau et al. (1999) paper presents the clumping factor approach in its essence, in the simplest possible flavor, and therefore is easy to reproduce.

Figure 11b now shows the most important quantity in the semi-analytical modeling: the porosity of the IGM $Q_{\rm H\,II}$ as a function of time. The dashed line shows the approximate expression from equation (23) of Madau et al. (1999) paper, and the solid line reflects a more accurate calculation based on equation (21) of Madau et al. (1999) paper with the comoving emissivity extracted from the simulation. Again, the main conclusion of Madau et al. (1999) is confirmed, the overlap occurs when $Q_{\rm H\,II} \approx 1$.

I have just demonstrated, that my simulation agrees (at least on a semi-qualitative level) with semi-analytical models of Miralda-Escudé et al. (1999) and Madau et al. (1999). Then how about the disagreement between the two? Clearly, if both agree with the same simulation, they also agree with each other. Let me know discuss the "differences" between the two semi-analytical approaches.

First of all, if one looks carefully, one can notice that equation (23) of Madau et al. (1999) is indeed the ratio of ϵ to R using Miralda-Escudé et al. (1999) notation, and the Miralda-Escudé

et al. (1999) equation (3) $\epsilon \approx R$ is indeed equation (24) of Madau et al. (1999). In other words, the dashed line in Fig. 11b is a ratio of the two curves in Fig. 9a. There is therefore no principal difference between the two approaches, but the advantage of Miralda-Escudé et al. (1999) approach is that they propose a model for the clumping factor based on the realistic density distribution function, rather than make an ad hoc assumption about the time evolution of the clumping factor. The main disadvantage of their model is that it is not necessarily accurate quantitatively.[5]

There however still remains a question of the number of ionizing photons per baryon required to ionized the universe. It cannot be simultaneously one (Miralda-Escudé et al. 1999) and many (Madau et al. 1999), can it?

The irony of the situation is that it indeed can. The difference between these two approaches is mostly terminological, and it can be best understood if one considers the number of ionizing photons per baryon as a function of scale, $N_{\gamma/b}(R)$ (this quantity cannot be rigorously defined mathematically, but has an intuitively clear physical sense). If Madau et al. (1999) count all photons emitted by all sources (in our case stars), i.e. they consider a quantity $N_{\gamma/b}(0)$ (here zero actually means the stellar surface), then Miralda-Escudé et al. (1999) use $N_{\gamma/b}(\mathrm{IGM})$ where the symbol IGM stands for the photons that escape local absorption, i.e. absorption in the gas bound to a given source, and corresponds to a range of scales of the order of $0.1 - 1\,\mathrm{Mpc}$. The top panel of Fig. 10 gives yet another quantity, $N_{\gamma/b}(R_{\mathrm{min}})$, where R_{min} is the resolution of the simulation, and in my case it is

$$R_{\mathrm{min}} = \frac{1}{1+z} h^{-1}\,\mathrm{kpc}$$

in physical units. It is now easy to bridge a difference between Madau et al. (1999) and Miralda-Escudé et al. (1999). The number $N_{\gamma/b}(0)$ is perhaps a factor of a few larger than $N_{\gamma/b}(R_{\mathrm{min}}) \approx 10$ at $z = 7$, which gives a value for $1/N_{\gamma/b}(0)$ of the order of a few percent, in complete agreement with Madau et al. (1999). On the other hand, as has been discussed above, $N_{\gamma/b}(\mathrm{IGM})$ should be considerably smaller than $N_{\gamma/b}(R_{\mathrm{min}})$, which leaves us with the conclusion of Miralda-Escudé et al. (1999) that $N_{\gamma/b}(\mathrm{IGM})$ is of the order unity.

There remains one final sticking point: a question of whether reionization is "fast" or "slow". But I have already mentioned in §3.1 that it is to a some degree a question of terminology as well. The whole process of reionization consists of three different stages: the "slow" (of the order of the Hubble time) pre-overlap, the "fast" (of the order of one tenth of the Hubble time) overlap, and the "slow" post-overlap. Miralda-Escudé et al. (1999) focused on the post-overlap stage, but considered the whole process of reionization, and therefore concluded that reionization was slow. [6] Madau et al. (1999) (and many other authors including Gnedin & Ostriker 1997) used the term

[5]Miralda-Escudé et al. (1999) model will work much better quantitatively for the case when the universe was reionized by a few bright quasars. In this case the amount of gas in the vicinity of sources is much smaller than in the case of stellar sources of reionization, when there are sources in a good fraction of all high density regions.

[6]There may however still remain a difference, since Miralda-Escudé et al. (1999) also claimed that the overlap is

"reionization" to label the overlap stage only, and thus concluded that reionization was fast. Thus, all the semi-analytical models give more-or-less similar predictions for the physics of reionization, and these predictions are reproduced by numerical simulations.

3.5. Reionization and the Lyman-alpha forest

While a question of reionization has its own astrophysical interest, reionization also leaves an imprint on the IGM at lower redshift, i.e. on the Lyman-alpha forest. One of the important properties of the forest is the so-called "effective equation of state", i.e. a tight relationship between the gas density and temperature in the low density regime, which is usually can be approximated by a power-law over a limited range of densities around the cosmic mean,

$$T \approx T_0(1+\delta)^{\gamma-1}$$

(Hui & Gnedin 1998). The time evolution of two parameters, T_0 and γ, depends on the thermal history of the universe, and therefore is tightly coupled to the processes taking place during reionization.

Usually, the evolution of the equation of state is calculated using the optically thin approximation and assuming that the radiation field is uniform (Hui & Gnedin 1998). This is clearly an approximation which needs to be verified.

Figure 12 shows the comparison of the "effective equation of state" from the simulation and the one computed in the optically thin approximation using the evolution of the spatially averaged radiation field extracted from the simulation (and thus both calculations have precisely the same mean J_ν as a function of time). While the two calculations are quite different at high redshift, the optically thin approximation gives a reasonably accurate answer at $z \lesssim 5$.

It is however important to notice, that the relationship between the gas density and temperature is more complicated at around the epoch of overlap and before. Figure 13 shows the joint mass-weighted distributions of the gas density and temperature at four redshifts, corresponding to Fig. 5 and Fig. 6. One can see that the power-law fits to the "effective equation of state" give poor representation of the true temperature-density relation at $z \gtrsim 6$.

Another important consideration that severely affects our ability to model the Lyman-alpha forest is how homogeneous the ionizing background is. Figure 14 shows the mean and rms fluctuation in the local ionizing background as a function of gas density at three values of redshift. As can be seen, there is essentially no evolution in the spatial distribution of the ionizing radiation

---

"gradual". It is not clear however without further elaboration of the Miralda-Escudé et al. (1999) model whether this "gradual" overlap is actually "fast" enough to be compatible with other semi-analytical approaches and with this work (after all, the ionizing intensity has to increase by two to three orders of magnitude over a Hubble time or so), or whether there is a genuine disagreement between Miralda-Escudé et al. (1999) and this work.

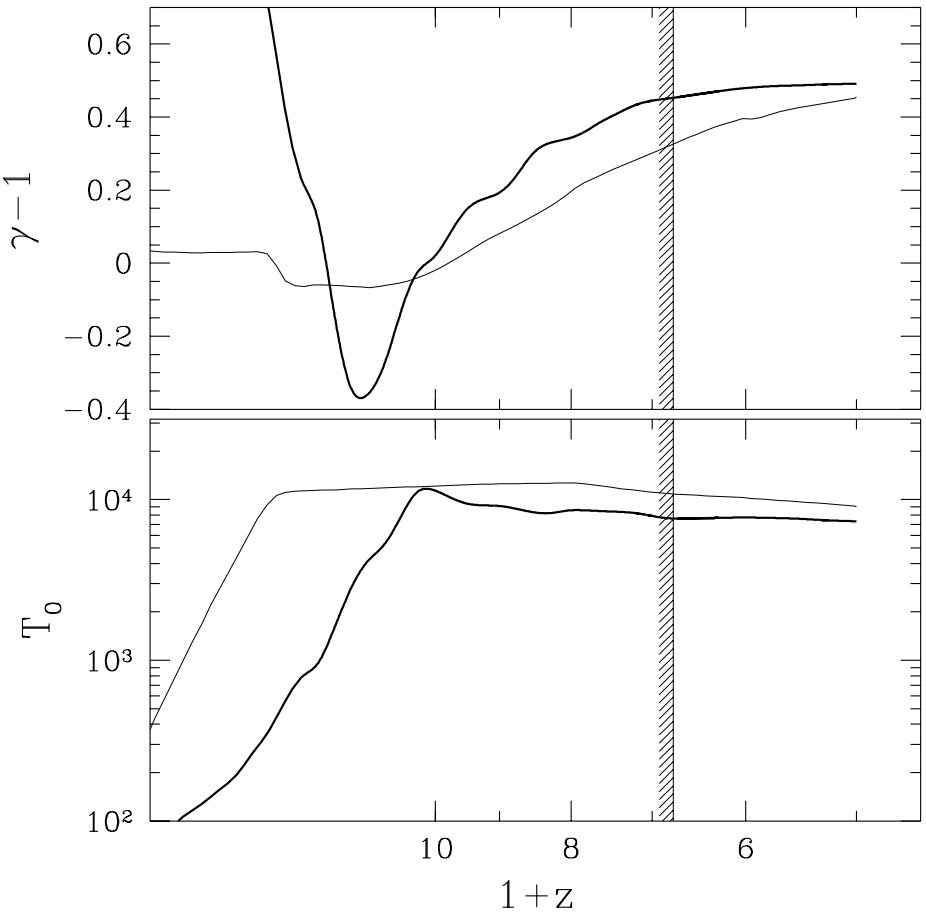

Fig. 12.— Evolution of the "effective equation of state" of the IGM, as represented by parameters T_0 (*bottom panel*) and γ (*top panel*), for the full numerical simulation (*bold line*) and the optically thin approximation (*thin line*). The shaded line marks the redshift above which the temperature-density correlation is not well established.

in the low density regime, and gas with overdensities in excess of about 10 is subject to severe (more than 30%) inhomogeneities in the ionizing background.

Since in the low density regime there exist a tight correlation between the column density of an absorption line and the gas density the absorption line originates in, it is possible to convert the gas density into the column density of the absorption line. Assuming $J_{21} = 0.5$ independently of redshift, and using equation (18) of Ricotti, Gnedin, & Shull (1999), I obtain:

$$N_{\rm H\,I} \approx 1.13 \times 10^{13}\,{\rm cm}^{-2}(1+\delta)^{1.52} \quad \text{at } z = 4,$$

and

$$N_{\rm H\,I} \approx 2.57 \times 10^{12}\,{\rm cm}^{-2}(1+\delta)^{1.66} \quad \text{at } z = 3.$$

Thus, if a 30% fluctuation in the ionizing background is still compatible with the background being

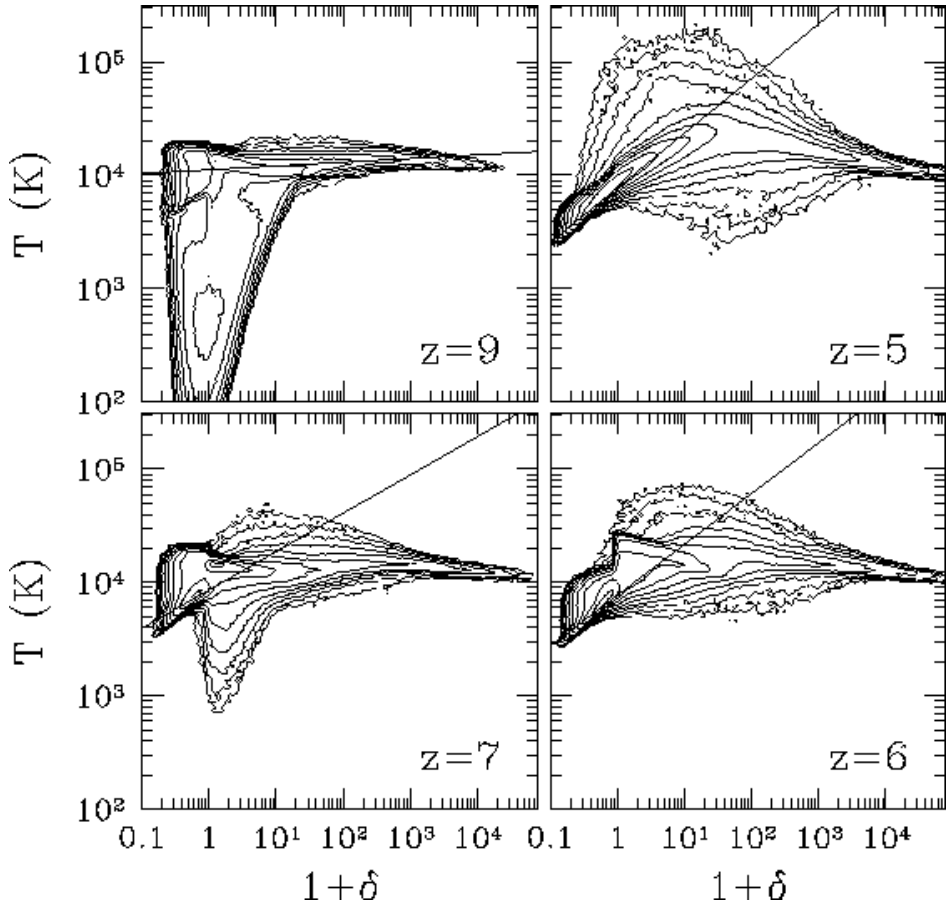

Fig. 13.— The joint mass-weighted distribution of the gas density and temperature in the production run N128_L4_A at four different redshifts. Bold solid lines show the power-law fits to the "effective equation of state" at the respective redshifts.

uniform, only the Lyman-alpha forest with $N_{\rm H\,I} < 4 \times 10^{14}\,{\rm cm}^{-2}$ at $z = 4$ and $N_{\rm H\,I} < 1 \times 10^{14}\,{\rm cm}^{-2}$ at $z = 3$ can be modeled in a uniform ionizing background approximation. If an accuracy of 10% is required (which corresponds to $\delta < 1$), then these limits shrink further to $3 \times 10^{13}\,{\rm cm}^{-2}$ and $8 \times 10^{12}\,{\rm cm}^{-2}$ respectively.

It is important to remember, that the production run simulation presented in this paper has still not converged numerically to below several tens of percent accuracy, and thus the numbers quoted above should be considered as only "within-a-factor-of-two" limits. Thus, it is likely to be safe to claim that theoretical models that assume a homogeneous ionizing background cannot reproduce the Lyman-alpha forest with the column density

$$N_{\rm H\,I} > 10^{15}\,{\rm cm}^{-2} \left(\frac{1+z}{5}\right)^6$$

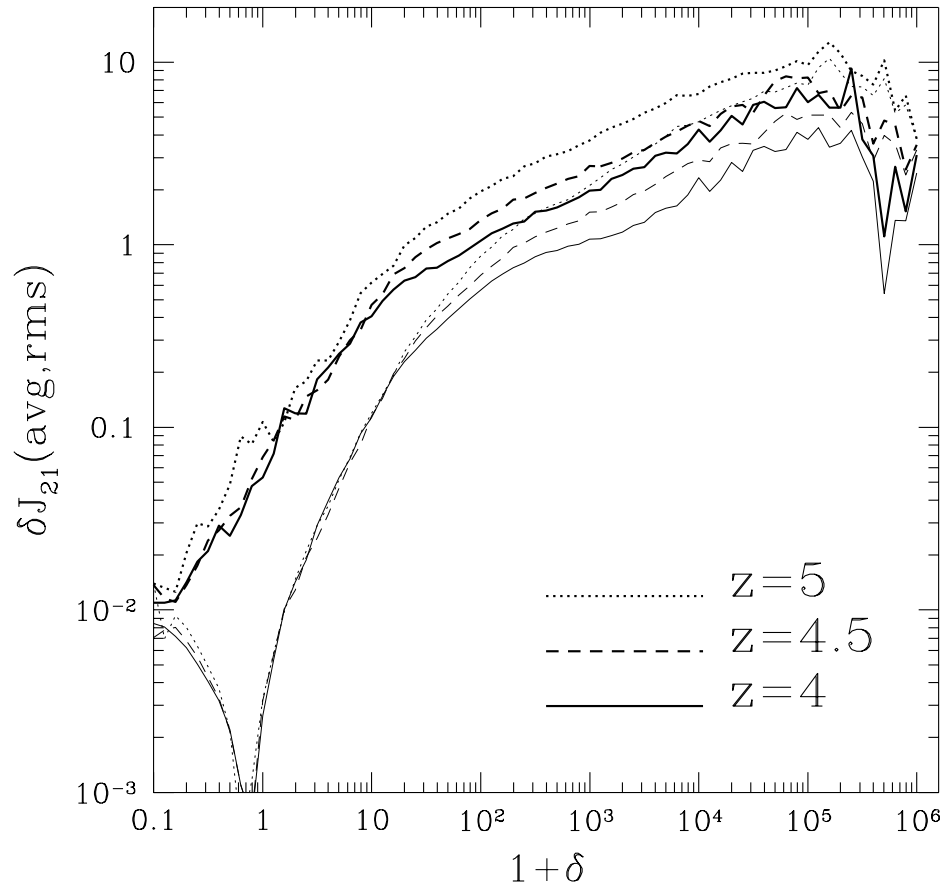

Fig. 14.— Average (*thin lines*) and rms (*bold lines*) fluctuation in the ionizing background as a function of gas density at $z = 5$ (*dotted lines*), $z = 4.5$ (*dashed lines*), and $z = 4$ (*solid lines*).

at a better than 30% level, and the forest with

$$N_{\mathrm{H\,I}} > 10^{14}\,\mathrm{cm}^{-2}\left(\frac{1+z}{5}\right)^6$$

with a better than 10% accuracy for z between 3 and 4 on a line-by-line basis.[7]

4. Conclusions

Cosmological numerical simulations that incorporate the effects of radiative transfer show that the whole process of reionization can be separated into three main stages: on a "pre-overlap" stage, which occupies a Hubble time or so, the H II regions expand into the neutral low density

[7]Of course, these errors may average out for some of the global properties of the forest.

IGM, leaving high density protrusions behind. The high density regions in the vicinity of a source are being ionized as well, albeit with a much slower rate. Pre-overlap culminates with the overlap, when the H II regions finally merge and the universe becomes transparent on a much larger scale than during the previous stage. The overlap epoch is characterized by a sharp rise (in about 10% of the Hubble time) in the level of the ionizing background and in the mean free path. At the "post-overlap" stage the universe consists of highly ionized low density gas, and still neutral high density gas, which is being gradually ionized as more and more ionizing photons become available. The boundary between the neutral and the ionized gas moves toward higher densities with time.

In a currently fashionable CDM-type cosmological scenarios the whole process of reionization occupies a considerable part of the early evolution of the universe from $z \sim 15$ until $z \sim 5$.

The aftermath of reionization has a profound effect on the IGM at lower redshift: the residual fluctuations in the ionizing background become quite significant at the column densities in the range from $10^{14}\,\mathrm{cm}^{-2}$ to $10^{15}\,\mathrm{cm}^{-2}$ (at $z = 3 - 4$) and may compromise theoretical models that do not take these fluctuations into account.

I am grateful to Jordi Miralda-Escudé and the referee Tom Abel for valuable comments that significantly improved the original manuscript. This work was partially supported by National Computational Science Alliance under grant AST-960015N and utilized the SGI/CRAY Origin 2000 array at the National Center for Supercomputing Applications (NCSA).

A. Local Optical Depth Approximation

As was shown in Gnedin & Ostriker (1997, equation [B9]), the energy density of radiation per unit frequency $I_\nu(x^i, t)$ can be represented as

$$I_\nu(x^i,t) = \bar{I}_\nu(t) e^{\bar{\tau}_\nu(t) - \tau_\nu(x^i,t)} + \frac{a}{4\pi c} \int d^3x_1 \frac{S_\nu(x_1^i,t) - \bar{S}_\nu(t)}{(x^i - x_1^i)^2} e^{-\tau_\nu(x^i,x_1^i,t)}, \tag{A1}$$

where $\bar{I}_\nu(t)$ is the volume average of I_ν, S_ν is the source function, $\tau_\nu(x^i, x_1^i)$ is the optical depth between x^i and x_1^i, $\tau_\nu(x^i, t)$ is the optical depth from a given point x^i to an "average" point in the universe, obtained by solving the homogeneous radiative transfer equation with no source function, and $\bar{\tau}_\nu(t)$ is the normalization constant obtained from the following condition:

$$e^{-\bar{\tau}_\nu(t)} \equiv \langle e^{-\tau_\nu(x^i,t)} \rangle$$

with brackets denoting the volume average.

In order to implement a radiative transfer scheme into modern cosmological simulations, one needs to be able to solve equation (A1) in $O(N_B)$ operations (where N_B is the number of baryonic resolution elements, and I ignore a possible logarithmic multiplier, so that $O(N_B \log N_B)$ still counts as $\mathrm{O}(N_B)$). However, equation (A1) in general includes two operations that are

more expensive than this count: (1) the one-point optical depth $\tau_\nu(x^i, t)$ requires $O(N_B \times N_D)$ operations, where N_D is the number of directions, and in general is of order of $N_B^{2/3}$, and (2) direct evaluation of the integral over the source function requires $O(N_B \times N_S)$ operations, where N_S is the number of sources. In the case under consideration, N_S is the number of stellar particles in the simulations, and is typically comparable to the number of baryonic resolution elements (perhaps, a factor of a few smaller, but it is the scaling that matters). Thus, direct (i.e. exact) implementation of equation (A1) is not feasible at the moment. [8]

Thus, one has to use approximations to try to decrease the operation count in evaluating (A1). The "Local Optical Depth" approximation to evaluate the first term in (A1) was developed in Gnedin & Ostriker (1997) and is based on a simple observation that high precision in evaluating $\tau_\nu(x^i, t)$ is not required: if the optical depth is small, it does not matter much whether it is 10^{-3} or 10^{-5}. If it is large, it again does not matter much whether it is 10^3 or 10^5. Only when $\tau \sim 1$ it needs to be computed accurately, but the volume of space where this condition is achieved is very small, so at the end no large error is introduced even if τ is computed with only an order-of-magnitude accuracy. Based on this consideration, the first *ansatz* is introduced in the solution of equation (A1):

$$\tau_\nu(x^i) = \sum_\alpha \sigma_\nu^{(\alpha)} n^{(\alpha)}(x^i) L(x^i), \tag{A2}$$

where index α runs over the list of species (in our case H I, He I, and He II), and L is a characteristic length, which is taken to be the characteristic length of the density distribution, and I drop the time dependence (present in all terms) hereafter for the sake of simplified notation. Specifically, I adopt the following expression for L:

$$L = \frac{1}{\sqrt{\alpha \left|\nabla \log \rho\right|^2 + \beta\rho \left|\Delta \log \rho\right| + \gamma \left|\nabla \log x_{\rm H\,I}\right|^2}}, \tag{A3}$$

where ρ is the baryon density, $x_{\rm H\,I}$ is the neutral hydrogen fraction, $\alpha = 0.216$ and $\beta = 0.068$ are constants [9] chosen to reproduce the correct result for the column density of a $1/(r^2 + r_c^2)$ density distribution in two limits $r \to 0$ and $r \to \infty$ and the last term is designed to work only for uniform density distribution, when L otherwise would be infinite. To this end I choose

$$\gamma = 30 \max(0, 1 - \rho/\bar{\rho}), \tag{A4}$$

where $\bar{\rho}$ is the average density of the universe and the factor 30 in front is chosen in order to satisfy the test D discussed above. Physically, this coefficient means that the approximation described in this paper spreads ionization fronts on average by a factor of $\sqrt{30} \sim 5$.

---

[8] In the case of reionization by quasars, however, when the number of sources N_S is small enough, the integral in (A1) can indeed be computed directly in $O(N_B)$ operations (Abel et al 1999).

[9] Note, that in Gnedin & Ostriker (1997) these constants are listed incorrectly because of a typo; they had proper values in the simulations described in that paper.

The Local Optical Depth approximation ensures that the first term in equation (A1) is computed in $O(N_B)$ operations.

Let me know turn to the second term. In general, integrals similar to the one in (A1) can only be computed in $O(N_B)$ operations with high quasi-Lagrangian resolution if they are convolutions. Thus, one of the ways to make this integral computable is to approximate it as a sum of several convolutions. For this purpose I divide all sources into two categories, "far" and "near" sources, and rewrite equation (A1) as

$$I_\nu(x^i) = \bar{I}_\nu e^{\bar{\tau}_\nu - \tau_\nu(x^i)} + I^F_\nu(x^i) + I^N_\nu(x^i). \tag{A5}$$

For a fluid element in a cosmological simulation, a source can be considered to be "far" if it sits in a separate clump. Then, the optical depth between the fluid element and the source can be approximated as a sum of one-point optical depths from two points:

$$\tau_\nu(x^i, x^i_1) \approx \tau_\nu(x^i) + \tau_\nu(x^i_1),$$

and, thus,

$$I^F_\nu(x^i) = \frac{a}{4\pi c} e^{-\tau_\nu(x^i)} \int d^3x_1 \frac{\left[S_\nu(x^i_1) - \bar{S}_\nu\right] e^{-\tau_\nu(x^i_1)}}{(x^i - x^i_1)^2}. \tag{A6}$$

This integral is a convolution, and can be computed in $O(N_B)$ operations by means of a standard P^{3}M technique. More than that, since the SLH-P^{3}M code already incorporates a P^{3}M gravity solver, it has been straightforward to modify the existing solver to compute the (A1) term.

The second, "near" term, $I^N_\nu(x^i)$, can be considered to include everything that is left to include in equation (A5) after the approximation (A6) is incorporated. However, in order to overcome the $O(N_B \times N_S)$ count, I must introduce another ansatz. For a "near" source, the two-point optical depth $\tau_\nu(x^i, x^i_1)$ can be decomposed into Taylor series,

$$\tau_\nu(x^i, x^i_1) \approx \nabla_j \tau_\nu(x^i, x^i)(x^j - x^j_1).$$

However, even this is not enough since the factor $\nabla_i \tau_\nu(x^i, x^i)$ still prevents the integral in (A1) from being a convolution. Thus, I introduce the second *ansatz* in the following form:

$$I^N_\nu(x^i) = \frac{a}{4\pi c} \int d^3x_1 \frac{S_\nu(x^i_1) - \bar{S}_\nu}{(x^i - x^i_1)^2} e^{-S|x^i - x^i_1|} \left[1 - e^{-\tau_\nu(x^i_1) - \tau_\nu(x^i)}\right]. \tag{A7}$$

Here S is a constant, independent of position (but it still may depend on time), and the integral now is a sum of two convolutions. The square bracket ensures that in the optically thin limit, $\tau \to 0$, the "near" term vanishes and the "far" term reduces to

$$I^F_\nu(x^i) = \frac{a}{4\pi c} \int d^3x_1 \frac{S_\nu(x^i_1) - \bar{S}_\nu}{(x^i - x^i_1)^2},$$

which is indeed the integral in equation (A1) in the limit $\tau \to 0$. Thus, the proposed approximation *becomes exact* in the optically thin limit.

The fact that the factor S is the same for all sources implies that the ionization front speed is correct only "on average", and each particular ionization front may propagate with a speed different from the correct one, but they will merge at approximately the right moment. Of course, if there is only one ionization front in the simulation, its speed is computed correctly.

From the argument presented above, it is clear that a good form for the factor S would be the following one:

$$S(t) = C(t)\langle\frac{\tau_{HI}(x^i,t)}{L(x^i,t)}\rangle, \tag{A8}$$

where $\tau_{HI}(x^i)$ is the one-point optical depth at the hydrogen ionization threshold, and the ratio τ_{HI}/L represents the gradient of the optical depth with the desirable feature that it is independent of the poorly defined quantity L.

The quantity C, which in part determines the speed of ionization fronts, needs to be fit to the tests. In other words, there is noting in the approximation (A5-A8) which ensures photon number conservation, so the quantity C should be chosen so that this number is at least approximately conserved.

Finally, a few words are in order about the frequency dependence. For stellar sources, considered in this paper, the shape of the source function is the same everywhere, so I can take it out of the integrals,

$$S_\nu(x^i,t) = g_\nu\rho_{MS}(x^i,t),$$

where g_ν is constant in space and time, and ρ_{MS} is the mass density of massive stars. However, this still leaves the frequency dependence of the optical depth in I^F and I^N. With the current computer capabilities it presents a considerable expense to perform the calculation of the integral in I^N and I^F at a sufficient number of frequency values.[10] Therefore, I introduce the third *ansatz* in the described approximate scheme. I introduce three effective column densities $N_{\rm eff}^{(\alpha)}$, where α again runs over H I, He I, and He II, so that

$$I_\nu^F = I_{\nu=0}^F \exp\left(-\sum_\alpha N_{\rm eff}^{(\alpha)}\sigma_\nu^{(\alpha)}\right),$$

and analogously for I^N. Then the integrals in I^F and I^N only need to be computed at zero frequency and at three threshold frequencies of H I, He I, and He II. In the future, when computer power increases by another factor of 10 or so, it will be possible to get rid of the third ansatz and compute integrals in I^F and I^N over a large number of frequency bins.

[10] Tests show that uniform logarithmic sampling of the frequency space should be at least 50 points per decade, which makes the total number of required frequency bins at least 200.

B. Testing the Approximation

Obviously, the proposed approximation is so simple that it is highly unlikely that it will work in all circumstances. Thus, my goal is to make sure that it works in cosmologically relevant conditions. This can be achieved by appropriately "training" the scheme against spherically symmetric numerical solutions, which can be obtained from direct integration of equation (A1).

For the test, I choose the following density distribution:

$$\rho(r) = \frac{\rho_c r_c^n}{(r^3 + r_c^3)^{n/3}}. \tag{B1}$$

This density distribution peaks at the center with $\rho = \rho_c$, has a core radius of r_c (taken to be 11 comoving kiloparsecs in all cases), and falls off as r^{-n} outside of the core radius. The specific form of the density law was chosen to make the coordinate transformation from real space x^i to Lagrangian space q^i,

$$\rho d^3x = \bar{\rho} d^3q,$$

analytical,

$$r = \left[\left(\frac{(3-n)\bar{\rho} q^3}{3\rho_c r_c^n} + r_c^{3-n} \right)^{3/(3-n)} - r_c^3 \right]^{1/3}.$$

This distribution is then embedded into the uniform mesh $x^i = q^i$ at a point r_m where the average density inside the sphere with radius r_m is $\bar{\rho}$, so that $\langle\rho\rangle = \bar{\rho}$ over the computational box. Tests are then performed with the 64^3 computational Lagrangian mesh and compared to the results of a spherically symmetric calculation. Spherically symmetric calculations are done with the same density profile and with 4000 shells between 0.1 kpc to 100 kpc. It has been verified that this number of shells is sufficient to achieve a 1% accuracy in the exact solution.

As the first test, I adopt a $n = 2$ ($\rho \propto r^{-2}$ ar large radii) density distribution with $\rho_c/\bar{\rho} = 10^3$ (maximum overdensity of 1000) embedded in the uniform radiation field with $\Gamma = 10^{-12}\,\mathrm{s}^{-1}$ and $\Gamma = 10^{-14}\,\mathrm{s}^{-1}$ (which roughly correspond to the radiation field after and before the overlap) at $z = 9$. No point source is included in this test, so it designed to test only the expression for the optical depth (A2,A3). Figure 15 now shows the temperature and the ionization fraction for the exact and the approximate solutions taken at some arbitrary moments in time for the two values of the photoionization rate. As one can see, the neutral hydrogen profiles are reproduced reasonably well, except just after the ionization front, whereas temperature profiles are reproduced less accurate, with extra heating observed just after the ionization front. As I will show below, the same behavior is observed in the test cases with the source places at the center and the ionization front propagating outward. This feature has to be considered an error of the LOD approximation and cannot be easily repaired.

Next, I perform a series of test with a source of ionizing radiation placed at the center of the spherically symmetric density distribution. The strength of the source is parametrized by the photoionization rate at the core radius Γ_c computed in the optically thin regime.

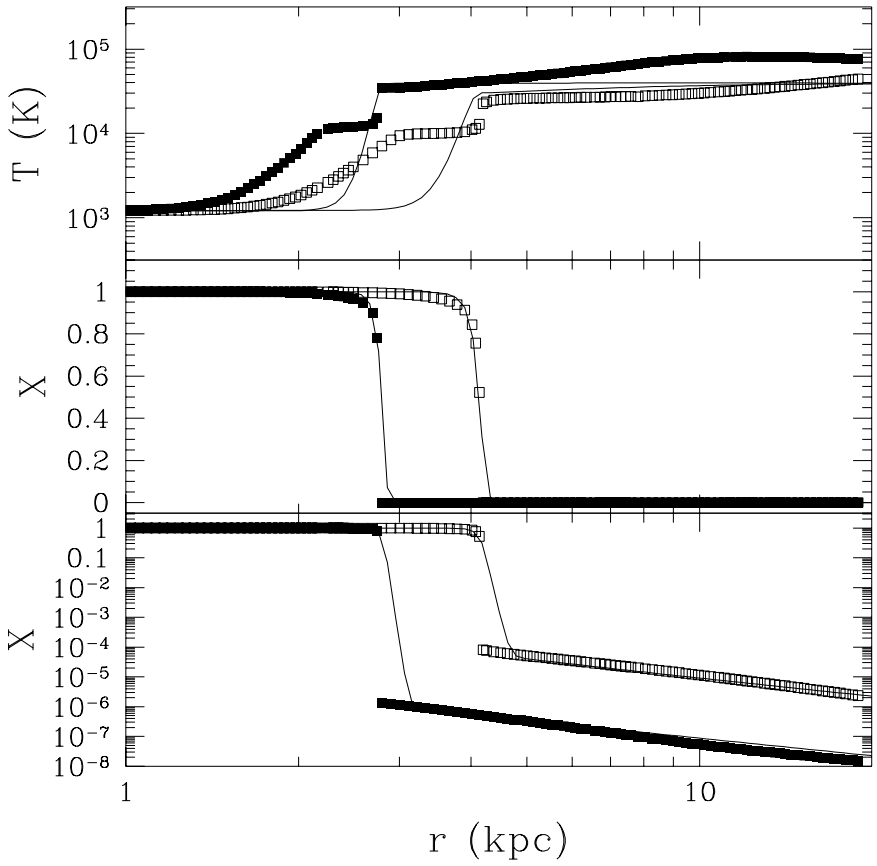

Fig. 15.— Distribution of the temperature (*top panel*) and neutral hydrogen fraction on linear (*middle panel*) and logarithmic (*bottom panel*) scale for a homogeneous radiation field test shining on the $n = 2$ and $\rho_c/\bar{\rho} = 10^3$ density distribution from outside at $z = 9$ with the photoionization rates of $\Gamma = 10^{-12}\,\mathrm{s}^{-1}$ and $\Gamma = 10^{-14}\,\mathrm{s}^{-1}$ taken at an arbitrary moments in time. Solid lines show the exact spherically symmetric solutions, and symbols mark the approximate 3D solutions, with solid squares corresponding to $\Gamma = 10^{-12}\,\mathrm{s}^{-1}$ and open squares corresponding to $\Gamma = 10^{-14}\,\mathrm{s}^{-1}$.

The seven main tests performed are listed in Table 2. The first five tests include the density distribution that falls off as r^{-2} at large radii, the test F verifies the propagation of the ionization front over the uniform density field ($\rho_c/\bar{\rho} = 1$), and the last test G includes the r^{-1} density law as the intermediate case between the tests A-E and F. That test is not used to train the approximate scheme, but rather to verify the performance of the approximation in a substantially different arrangement.

Additional tests have been performed to verify numerical convergence and other technical issues such as dependence on the cell size.

As the results of fitting the approximate scheme to the exact solutions of the first six tests,

Table 2. Parameters of Test Cases

Case	z	$\Gamma_c a^3$ (s^{-1})	$\rho_c/\bar{\rho}$	n	t_0 (yr)
A	4	10^{-14}	10^3	2	1×10^6
B	4	10^{-12}	10^3	2	2×10^4
C	4	10^{-10}	10^3	2	4×10^2
D	9	10^{-14}	10^3	2	1×10^6
E	9	10^{-12}	10^3	2	2×10^4
F	9	10^{-14}	1	–	1×10^6
G	9	10^{-12}	40	1	1.5×10^4

the following expression for the coefficient C is derived:

$$C = \frac{1}{0.1 + Q^{2/3}}, \tag{B2}$$

where

$$Q(t) = 8.6 \int_0^t \frac{\tilde{\Gamma}(t')}{a(t')^{1/3}} dt',$$

and

$$\tilde{\Gamma} = \max\left(0, \langle \Gamma - \frac{90}{C^{1/3}} R(T) n_e \rangle\right),$$

where Γ is the photoionization rate, R is the recombination coefficient, and n_e is the electron number density. The quantity C thus depends on the time integral of the volume averaged photoionization rate (with a correction for recombinations), and the factor $a^{1/3}$ in the denominator accounts for the redshift of ionizing photons.

Equation (B2) thus completes the approximation. I would like to stress here that equation (B2) is a purely empirical fit and may have no physical basis behind it. It appears physically plausible since it is based on the number of ionizing photons emitted over the course of action, but its specific form cannot be "deduced" from the first principles (in large part due to the fact that the Local Optical Depth approximation is based on two ansatzes).

We can now investigate how the developed approximation works in test cases. Figure 16 shows the comparison between the exact and approximate calculations of the evolution of the ionization front in all seven test cases. The agreement is good everywhere except cases D and E at early times, but this is due to a finite resolution of a numerical simulation. Deviations of the order of 20% are also observed in case A at late times. The case G shows in general a worse agreement, because this test was not used to fit the approximate scheme to the exact solutions. It can thus be used to demonstrate the level of accuracy of the LOD approximation for arrangements different

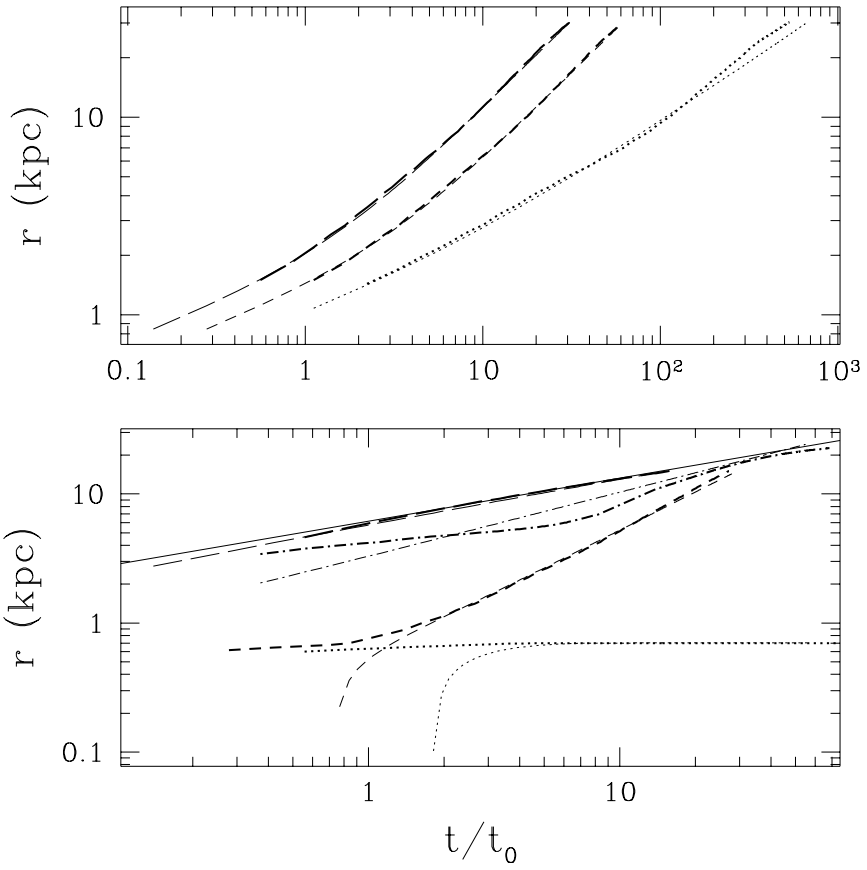

Fig. 16.— Evolution of the ionization front in seven test cases. The upper panel shows cases A (*dotted line*), B (*short-dashed line*), and C (*long-dashed line*). The lower panel shows cases cases D (*dotted line*), E (*short-dashed line*), F (*long-dashed line*), and G (*dot-dashed line*). Thin lines mark the exact solution, and bold lines show the Local Optical Depth approximation. Time axis is scaled arbitrarily, with the scaling factor listed in Table 2. The thin solid line shows the Shapiro analytical solution for the test F (homogeneous density field).

from those used to "train" the approximation. The solid thin line in Fig. 16 shows the analytical approximation of Shapiro (1986) and Shapiro & Giroux (1987) for the case F (the homogeneous density field). The small deviation observed at early times is due to the fact that the thickness of the ionization front in this regime is not negligible compared to the size of the H II region, and the Shapiro solution becomes invalid.

The question of numerical resolution in any 3D treatment of the radiative transfer is of utmost importance. Gravitational collapse proceeds from low density to high density, and therefore even with poor resolution low density structures can be reproduced reliably. The situation is just the opposite with the radiative transfer: the ionization front moves from high density regions into low density regions. Thus, during the first time-step the ionization front has to be fully resolved, or it will never leave the resolution element it originated in (unless it is specifically followed on

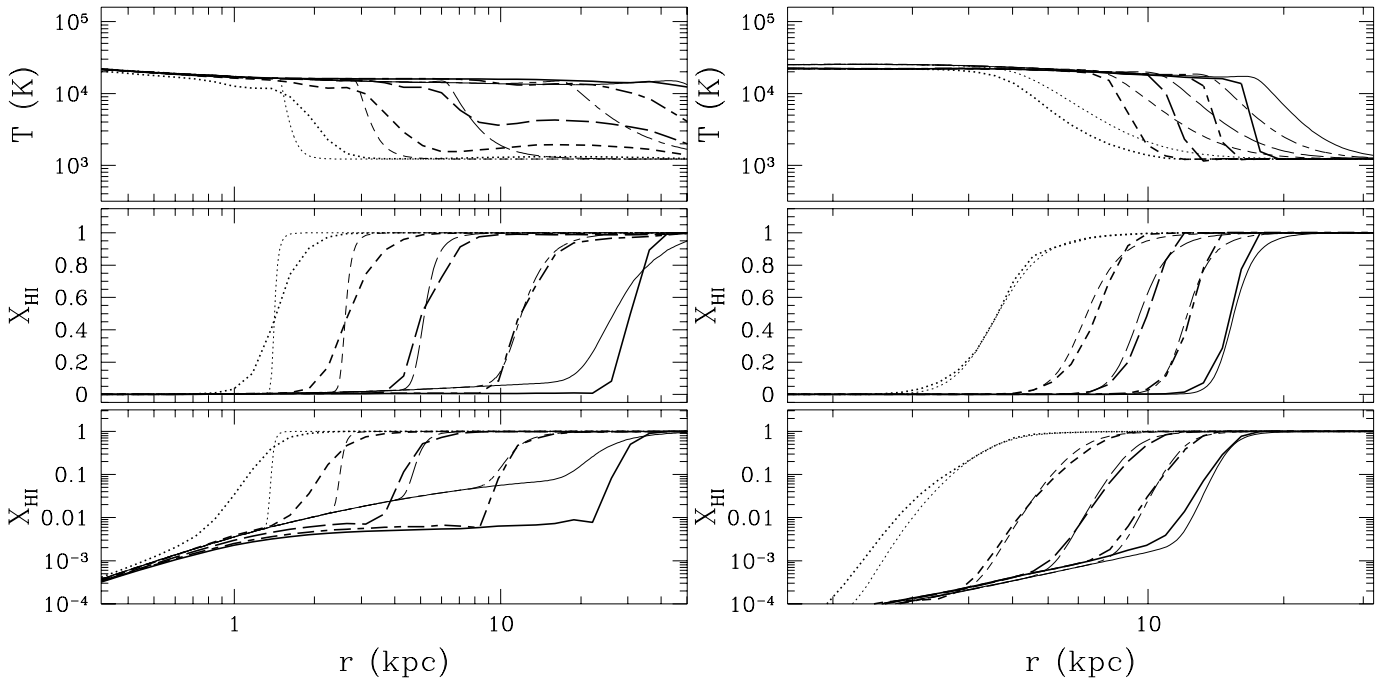

Fig. 17a Fig. 17b

Fig. 17.— Distribution of the temperature (*top panel*) and neutral hydrogen fraction on linear (*middle panel*) and logarithmic (*bottom panel*) scale at five different times (as labeled by different lines) in the exact solution (*thin lines*) and the approximate solution (*thick lines*) for cases A (panel *a*), D (panel *b*), and B (panel *c*).

inside that resolution element). Thus, to avoid the problem of ionization front getting "stuck" due to lack of resolution, I impose the condition that $S\epsilon > 2$, where ϵ is the gravitational softening length. In test cases E and F this leads to the ionization front deviating from the exact solution at earlier times, but catching up later, when it becomes fully resolved in the approximate scheme.

Figure 17 shows in three panels time-evolution of the temperature and the neutral hydrogen fraction for cases A, B, and D. Since the ionization front in the approximate solution is not sharp, I choose the location where hydrogen is 50% ionized as a location of the ionization front in the approximate solution for the purpose of producing Fig. 16.

Several defects of the approximate scheme are apparent in Fig. 17. First, the ionization front is spread over a considerable distance, but for cosmological purposes this distance is still much smaller than the correlation length, and thus this does not make the big impact on the dynamics of the gas. Second, at late times in Fig. 17a the approximate scheme significantly overestimates the degree of ionization, due to the fact that the optical depth between the position of the ionization front and the source is of the order of unity, and here the Local Optical Depth approximation is bound to fail. Third, in Fig. 17a the temperature ahead of the front is higher than it should be, whereas in Fig. 17b it is lower. This is due to the fact that leaking of high-energy photons

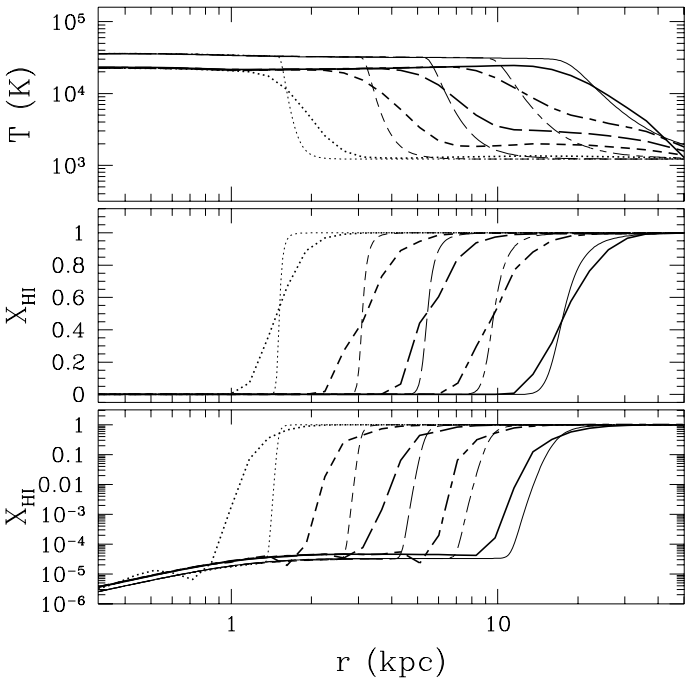

Fig. 17c

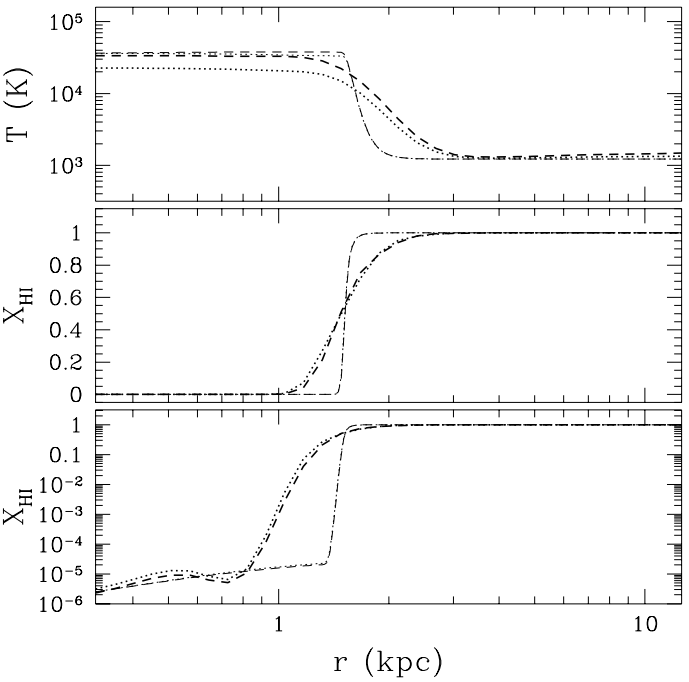

Fig. 18

Fig. 18.— Distribution of the temperature (*top panel*) and neutral hydrogen fraction on linear (*middle panel*) and logarithmic (*bottom panel*) scale in test B in the exact solution (*thin lines*) and the approximate solution (*thick lines*). Dotted lines show the full solution, and dashed lines show the solution with cooling disabled. Cooling along the ionization front leads to an incorrect prediction for the post-front temperature.

across the front is not captured correctly by the approximation. Finally, the most severe defect is apparent in Fig. 17c: the temperature inside the front in the approximate solution is some 50% lower than it should be. This is due to the fact that the ionization front is not sharp, and the gas is able to cool while being ionized, whereas in the exact solution the neutral fraction drops from 1 to 10^{-5} so quickly, that no appreciable cooling occurs. To illustrate this further I show in Figure 18 the first line from Fig. 17b together with the exact and approximate solution for the same test case but with disabled cooling. Without cooling, the approximate solution predicts the post-front temperature right, but with cooling it allows a considerable loss of energy.

This defect however is not specific to the Local Optical Depth approximation, but obviously will be present in any scheme that does not resolve the ionization front. It is possible to modify the SLH technique in such a way as to allow the mesh to deform appropriately and to achieve higher resolution at the ionization front, but this work is well beyond the scope of this paper.

Another significant defect of the Local Optical Depth approximation is the absence of shadowing. Since LOD approximation does not include proper ray tracing, the ionizing flux from a given source in a given fluid element is only influenced by the opacity in the vicinity of the source and the fluid element, and therefore a shadowing effect of a self-shielded neutral cloud on the line

of sight from the fluid element to the source would not be reproduced. Fortunately, shadowing is not important during cosmological reionization: a self-shielded cloud with the size of 10 kpc moving with the velocity of 100 km/ s produces a shadowing effect only during a time interval of 100 Myr, which is more than two orders of magnitude shorter than the recombination time in the low density IGM. Thus, the ionization and the thermal state of the IGM would not be affected by the passage of a cloud shadow.

REFERENCES

Abel, T., Norman, M. L., & Madau, P. 1999, ApJ, 523, 66

Bertschinger, E., & Gelb, J. 1991, Computers in Phys., 5, 154

Bond, J. R., & Jaffe, A. H. 1998, Phil. Trans. Royal Soc. Lon. A, in press (astro-ph 9809043)

Chiu, W. A., & Ostriker, J. P. 1999, in preparation

Ciardi, B., Ferrara, A. 1997, ApJ, 483, L5

Ciardi, B., Ferrara, A., Governato, F., & Jenkins, A. 1999, MNRAS, submitted (astro-ph 9907189)

Giallongo, E., D'Odorico, S., Fontana, A., McMahon, H. G., Savaglio, S., Cristiani, S., Molaro, P., & Treverse, D. 1994, ApJ, 425, L1

Giroux, M. L., & Shapiro, P. R. 19996, ApJS, 102, 191

Gnedin, N. Y. 1995, ApJS, 97, 231

Gnedin, N. Y., & Bertschinger, E. 1996, ApJ, 470, 115

Gnedin, N. Y., & Hui, L. 1998, MNRAS, 296, 44

Gnedin, N. Y., & Ostriker, J. P. 1997, ApJ, 486, 581

Griffiths, L. M., Barbosa, D., & Liddle, A. D. 1998, MNRAS, in press (astro-ph 9812125)

Haiman, Z., & Loeb, A. 1997, ApJ, 483, 21

Haiman, Z., & Loeb, A. 1998, ApJ, 503, 505

Hui, L., & Gnedin, N. Y. 1998, MNRAS, 292, 27

Lu, L., Sargent, W. L. W., Womble, D. S., & Takada-Hidai, M. 1996, ApJ, 472, 509

Madau, P. 1999, invited review at the VLT Opening Symposium, Antofagasta, Chile 1-4 March (astro-ph 9907268)

Madau, P., Haardt, F., & Rees, M. J. 1999, ApJ, 514, 648

Madau, P., Meiksin, A., & Rees, M. J. 1997, ApJ, 475, 429

Miralda-Escudé, J., Haehnelt, M., & Rees, M. J. 1999, ApJ, submitted (astro-ph/9812306)

Nagamine, K., Cen, R., & Ostriker, J. P. 1999, ApJ, submitted (astro-ph 9902372)

Ostriker, J. P., & Gnedin, N. Y. 1996, ApJ, 472, L63

Renzini, A. 1999, in "When and How do Bulges Form and Evolve?", eds. C.M. Carollo, H.C. Ferguson, R.F.G. Wyse (Cambridge: Cambridge University Press), in press (astro-ph 9902361)

Ricotti, M., Gnedin, N. Y., & Shull, J. M. 1999, ApJ, submitted (astro-ph 9906413)

Shapiro, P. R. 1986, PASP, 98, 1014

Shapiro, P. R., & Giroux, M. L. 1987, ApJ, 321, L107

Songaila, A., Hu, E. M., Cowie, L. L., & McMahon, R. G. 1999, ApJ, 525, L5

Steidel, C. C., Adelberger, K. L., Ciavalisco, M., Dickinson, M., & Pettini, M. 1999, to appear in the proceedings of the Xth Rencontres de Blois, "The Birth of Galaxies", July 1998 (astro-ph 9812167)

Tegmark, M., Silk, J., Rees, M. J., Blanchard, A., Abel, T., & Palla, F. 1997, ApJ, 474, 1

Valageas, P., & Silk, J. 1999, å, in press (astro-ph 9907068)

Williger, G. M. et al. 1994, ApJ, 453, L57